

Solving the Dual Problems of Dynamic Programs via Regression

Helin Zhu, Fan Ye and Enlu Zhou
School of Industrial and Systems Engineering,
Georgia Institute of Technology

Abstract

In recent years, information relaxation and duality in dynamic programs have been studied extensively, and the resulted primal-dual approach has become a powerful procedure in solving dynamic programs by providing lower-upper bounds on the optimal value function. Theoretically, with the so called value-based optimal dual penalty, the optimal value function could be recovered exactly via strong duality. However, in practice, obtaining tight dual bounds usually requires good approximations of the optimal dual penalty, which could be time-consuming due to the conditional expectations that need to be estimated via nested simulation. In this paper, we will develop a framework of regression approach to approximating the optimal dual penalty in a non-nested manner, by exploring the structure of the function space consisting of all feasible dual penalties. The resulted approximations maintain to be feasible dual penalties, and thus yield valid dual bounds on the optimal value function. We show that the proposed framework is computationally efficient, and the resulted dual penalties lead to numerically tractable dual problems. Finally, we apply the framework to a high-dimensional dynamic trading problem to demonstrate its effectiveness in solving the dual problems of complex dynamic programs.

Key words: Information relaxation, duality, dynamic program, optimal dual penalty, regression.

1 Introduction

Markov decision process (MDP) is a powerful model for complex dynamic decision making problems under uncertainty that arise in various fields, such as supply chain engineering, operations research, and financial engineering. The objective usually is to find an optimal policy that maximizes the expected accumulated rewards (or minimizes the expected accumulated costs) in the long run. These problems could be theoretically solved by the celebrated Bellman dynamic programming approach; however, in practice, it suffers from the so-called “curse of dimensionality”, meaning that the size of the state space and the action space, and hence the complexity of the program increases exponentially in the dimension of the problem. Therefore, it is rarely the case that the optimal policy of a real-world dynamic programming problem could be solved exactly. Facing this issue, abundant literature has been focusing on developing good approximate dynamic programming methods that aim to

construct good suboptimal policies, see Bertsekas (2007), Chang et al. (2007) and Powell (2011), etc. In principle, given a policy (usually suboptimal), Monte Carlo simulation could be used to evaluate the policy and generate good lower bound estimators on the optimal value function by simulating a large number of state-action sample paths under the policy. It scales well with the dimension of the underlying system. However, in lack of the exact optimal value function or its upper bounds, the quality of the policy and the optimality gap of the lower bounds are difficult to measure.

The duality theory developed independently by Rogers (2007) and Brown et al. (2010) addresses this issue by formulating and solving the dual representation of the (primal) dynamic program (DP), and providing upper bounds on the optimal value function. If the duality gap, i.e., the difference between the lower bound induced by the policy and the upper bound, is small enough, then one could claim that the policy is sufficiently good. The main idea of this duality theory is to relax the non-anticipativity constraint on all the feasible policies of the DP, i.e., allow the decision maker (DM) to choose actions based on the outcomes of future uncertainties, and penalize the DM for the access to the future information. Thus, this framework is also termed as information relaxation duality theory. In practical implementation, generating an upper bound on the optimal value function using the dual formulation only requires solving multiple pathwise deterministic inner optimization problems, and scales well with the dimension of the system via Monte Carlo simulation.

The information relaxation duality theory for general DPs originates from the dual theory in pricing American-style options, developed independently by Rogers (2002), Haugh and Kogan (2004), and Anderson and Broadie (2004), and further extended by Chen and Glasserman (2007). They are able to generate upper bounds on the option price by solving the associated dual problem, which is obtained by relaxing the non-anticipativity constraint on all the feasible exercising strategies (which essentially are stopping times) and penalizing the payoff function with a martingale adapted to the nature information filtration. Furthermore, if the penalizing martingale is the Doob-Meyer martingale component of the option price process, namely the “optimal dual martingale”, then strong duality is achieved, meaning that the upper bound is tight.

Rogers (2007) and Brown et al. (2010) generalize the idea of duality theory from American option pricing to general discrete-time DPs and provide a broader interpretation of the dual martingale. From Brown et al. (2010)’s perspective, the dual martingale could be regarded as the penalty for the access to the outcomes of future asset prices. Furthermore, similar to the existence of optimal dual martingale in option pricing, the authors show that in general DPs there also exist optimal dual penalties such that the dual problem with those penalties recovers the optimal value function. In particular, one specific choice of optimal dual penalties is the Doob-Meyer martingale component of the optimal value function process with respect to (w.r.t.) the natural information filtration at terminal horizon (referred to as the value-based optimal dual penalty). Following this line of research, there have been many new methodologies and applications in recent years. Brown and Smith (2011) and Brown and Smith (2014) propose the gradient/subgradient penalties for the dual problems of convex DPs (concave reward function and convex action sets), in order to preserve the convexity structure in the dual problems and facilitate the optimization. They further show that these penalties could be optimal under appropriate conditions. Ye and Zhou (2015)

generalizes the duality theory to controlled Markov diffusion (CMD) under a continuous-time setting and reveal the structure of the optimal dual penalty as a stochastic integral. Applications include but not limited to the valuation of gas storage Lai et al. (2010), performance measurements of trading strategies in portfolio management Moallemi and Saglam (2015), and robust multi-armed bandit problems Kim and Lim (2015).

We also note that there are many new developments in computational methods that aim at constructing good dual penalties. In general, the (value-based) optimal dual penalty could not be computed exactly because it involves the optimal value functions that are not available and conditional expectations that need to be estimated. The naive approach is to replace the optimal value functions with approximate ones, and use nested simulation to estimate the conditional expectations; however, this approach often requires substantial computational effort and might cause the resulted approximation to lose the dual feasibility. Various methods have been proposed to improve the accuracy and efficiency of the approximation, including the non-nested simulation approach by Belomestny et al. (2009) and Zhu et al. (2015) in American-style option pricing, and the pathwise optimization techniques by Desai et al. (2011) and Ye and Zhou (2012). The advantage of the pathwise optimization method is that it explores a subspace of feasible dual penalties by considering the best linear combination of the existing dual penalties. However, the drawback is that it requires solving a new stochastic optimization problem, which could be computationally expensive or intractable. Moreover, the quality of the resulted dual penalty heavily relies on the quality of the existing dual penalties; hence, good performance could not be guaranteed.

We notice that two key things are missing in most of the existing approaches. First, the structure of the optimal dual penalty is not well-studied. Since the space consisting of all feasible dual penalties is a function space (referred to as the dual penalty space hereafter), the optimal dual penalty could be viewed as a point in that space. Therefore, we could approximate it by computing and estimating all its coordinates w.r.t. a functional basis of the dual penalty space. If such estimation could be achieved without Monte Carlo simulation, then nested simulation is circumvented. Second, the dual problem is usually solved independently of the primal problem, meaning that some useful information (e.g., the suboptimal policy has been evaluated, and the simulations have been carried out) in the primal problem is not well-utilized. Properly utilizing that information might help facilitate estimation of the coordinates.

Motivated by these observations, in this paper we propose a framework of regression approach that explores the structure of the dual penalty space as well as the optimal dual penalty while utilizing the information from the primal problem. In particular, we find two function bases of the dual penalty space that possess desired properties such that the corresponding coordinates of the optimal dual penalty are easy to compute and the resulted approximations of the optimal dual penalty remain dual feasible. Furthermore, instead of using Monte Carlo simulation to estimate those coordinates, we propose an efficient regression method that reuses the suboptimal policy and the samples generated in the primal problem.

The proposed framework of regression approach has the following advantages. First, it is a scheme that requires minimal extra simulation or computational costs since it reuses the suboptimal policy and samples generated in the primal problem. Second, it generates

approximations of the optimal dual penalty without nested simulation suffered by some existing approaches. Therefore, computational efficiency is significantly improved. Third, in contrast to some of the existing approaches, the resulted dual penalties are feasible, and thus induce valid upper bounds on the optimal value function. Finally, as we will see later, the framework is robust. Dual penalties with desired structural properties could be generated by constructing proper regressors in the regression. For instance, for a convex DP, the framework is able to generate dual penalties that preserve convexity in the dual problem.

It turns out, in a broader sense, several existing approaches to approximating the optimal dual penalty could be regarded as special cases of this framework under specific settings. For instance, Belomestny et al. (2009) proposes a non-nested simulation approach to approximating the optimal dual martingale in American option pricing. In particular, the authors apply martingale representation theorem on the optimal dual martingale, then approximate the resulted stochastic integral with an Ito sum, and finally estimate the integrands in the Ito sum via regression. We will show that this approach could be viewed as a special case of the proposed framework with a specific functional basis of the dual penalty space. Our framework is more universal and powerful, because it reveals the structure of the optimal dual penalty regardless of the underlying probability measure (i.e., not restricted to the Brownian measure in Belomestny et al. (2009) or the Poisson random measure in Zhu et al. (2015)). We will also show that the approximation scheme in Ye and Zhou (2015) could be viewed as a special case of the proposed framework as well. To summarize, the contributions are as follows:

- We develop a framework of regression approach to approximating the optimal dual penalty in general DPs that (1) circumvents nested simulation that is suffered by some existing approaches, (2) requires minimal extra simulation or computational costs by reusing the suboptimal policy and samples generated in the primal problem, and (3) yields feasible dual penalties and valid upper bounds. Furthermore, the framework is quite robust in the sense that it is capable of generating dual penalties with desired structural properties by properly constructing the regressors in the regression.
- We explore the dual penalty space and reveal the structure of the optimal dual penalty for general DPs, which enable us to generate good feasible dual penalties in a systematic way. Several existing approaches could be regarded as special cases of the framework under specific settings.
- The application of the proposed framework to a high-dimensional dynamic trading problem shows that it is effective and efficient in solving the dual problems of complex DPs. In particular, it generates accurate approximations of the optimal dual penalty and tight upper bounds on the optimal value function when good suboptimal policies and appropriate regressors are used in the regression approach.

The rest of the paper is organized as follows. In Section 2, we review the basics of dynamic programming and information relaxation duality theory. We present the framework of regression approach in Section 3. In Section 4, we show several existing approaches are special cases of the proposed framework under specific settings. In Section 5, we apply the proposed framework to a high-dimensional dynamic trading problem, and demonstrate its effectiveness and efficiency in solving the dual problems of complex DPs. Conclusions are provided in Section 6.

2 Dynamic Programs and Dual Formulations

On a general probability space $(\Omega, \mathbb{P}, \mathcal{F})$, in which Ω is the set of all possible outcomes (scenarios) of uncertainties, \mathbb{P} is the underlying probability measure and \mathcal{F} is the σ -algebra consisting of all the events (measurable subsets of Ω). Consider a finite-horizon MDP as follows. Time is indexed by $\mathcal{T} = \{0, 1, \dots, N\}$. The state x follows the dynamics

$$x_{n+1} = f(x_n, a_n, z_{n+1}), \quad n = 0, 1, \dots, N-1, \quad (2.1)$$

where f is the deterministic transition function, $x_n \in \mathcal{X}_n$ denotes the state at period n that lives in the state space \mathcal{X}_n , $a_n \in \mathcal{A}_n$ denotes the action/control at period n that is chosen from the action space \mathcal{A}_n , z_{n+1} is the random noise at period n and $\{z_n : n = 1, \dots, N\}$ are assumed to be independently and identically distributed (i.i.d.) random variables with probability measure ρ on support space $\Xi \subseteq \mathbb{R}^d$. The evolution of information is described by the natural filtration $\mathcal{F} = \{\mathcal{F}_n : n = 0, \dots, N\}$. Loosely speaking, \mathcal{F}_n describes the information available to the DM at period n . In particular, each z_n is \mathcal{F}_n -measurable. Without loss of generality, we further assume $\mathcal{F}_0 = \{\emptyset, \Omega\}$, meaning that the DM initially does not have any knowledge about the outcomes of uncertainties, and $\mathcal{F}_N = \mathcal{F}$, meaning that the DM knows all the possible outcomes of uncertainties in the end.

We use a mapping $\alpha_n(\cdot)$ from the state space \mathcal{X}_n to the action space \mathcal{A}_n , i.e., $\alpha_n : \mathcal{X}_n \rightarrow \mathcal{A}_n$, to denote the decision rule at period n . A policy/strategy $\alpha := (\alpha_0, \alpha_1, \dots, \alpha_{N-1})$ consisting of a sequence of decision rules is called non-anticipative/ \mathcal{F} -adapted, if each decision rule $\alpha_n(\cdot)$ is \mathcal{F}_n -measurable. Intuitively, it means the DM chooses the action a_n only based on the information accumulated up to period n , he/she shall not choose the action a_n based on the future information. We use $\mathbb{A}_{\mathcal{F}}$ to denote the set of all non-anticipative policies, and \mathbb{A} to denote the set of all policies (including the anticipative ones); clearly $\mathbb{A}_{\mathcal{F}} \subseteq \mathbb{A}$. Furthermore, we associate an \mathcal{F}_n -measurable reward function $r_n(x_n, a_n)$ with the state dynamics to represent the immediate reward after the DM chooses action a_n at period $n = 0, \dots, N-1$, and an \mathcal{F}_N -measurable function $r_N(x_N)$ as the terminal reward at period N .

Given $x_0 \in \mathcal{X}_0$, the objective of the DM is to select a non-anticipative policy $\alpha \in \mathbb{A}_{\mathcal{F}}$ that maximizes the accumulated rewards over all the periods, i.e.,

$$(P) : V_0(x_0) \triangleq \sup_{\alpha \in \mathbb{A}_{\mathcal{F}}} \mathbb{E}_0 \left[\sum_{n=0}^{N-1} r_n(x_n, \alpha_n(x_n)) + r_N(x_N) \right], \quad (2.2)$$

where $\mathbb{E}_0[\cdot]$ means that the expectation is taken w.r.t. \mathcal{F}_0 , and we will use $\mathbb{E}_n[\cdot]$ to denote the expectation taken w.r.t. \mathcal{F}_n thereafter.

It is well known that problem (2.2) could be recursively solved theoretically via Bellman backward dynamic programming

$$\begin{cases} V_N(x_N) \triangleq r_N(x_N), \\ V_n(x_n) \triangleq \sup_{a_n \in \mathcal{A}_n} \{r_n(x_n, a_n) + \mathbb{E}_n[V_{n+1}(x_{n+1})]\}, \quad n = N-1, \dots, 0, \end{cases} \quad (2.3)$$

where $V_n(x_n)$ represents the optimal value function of the DP with initial period n and initial state x_n . However, in practice, the Bellman recursion (2.3) could hardly be solved

exactly for most cases, due to the curse of dimensionality. Therefore, one often needs to settle with suboptimal policies or approximate optimal value functions. A suboptimal policy could be evaluated conveniently via Monte Carlo simulation to generate a lower bound on the optimal value function. However, in the absence of the exact optimal value function or its upper bounds, the quality of the suboptimal policy could hardly be measured.

2.1 A Dual Formulation

The duality theory developed by Rogers (2007) and Brown et al. (2010) addresses the aforementioned issue by formulating a dual representation of the (primal) DP (2.2), and solving the dual problem provides a valid upper bound on the optimal value function. Therefore, the quality of a suboptimal policy could be empirically measured by examining the duality gap. If it is sufficiently small, then the suboptimal policy could be claimed to be near optimal.

To be more specific, let us rigorously define a *feasible dual penalty* as follows.

Definition 2.1. We say $M(\boldsymbol{\alpha}, \mathbf{z})$, a functional of policy $\boldsymbol{\alpha} \in \mathbb{A}$ and noise sequence $\mathbf{z} := (z_1, \dots, z_N)$, is a feasible dual penalty if

$$\mathbb{E}_0 [M(\boldsymbol{\alpha}, \mathbf{z})] = 0, \quad \forall \boldsymbol{\alpha} \in \mathbb{A}_{\mathcal{F}}. \quad (2.4)$$

Put in another way, a penalty function $M(\boldsymbol{\alpha}, \mathbf{z})$ is dual feasible if it does not penalize any non-anticipative policy in expectation. We further use $\mathbb{M}_{\mathcal{F}}$ to denote the set of all feasible dual penalties.

Remark 2.1. Definition 2.1 is slightly different from the one in Brown et al. (2010), in which a dual penalty is called feasible if $\mathbb{E}_0 [M(\boldsymbol{\alpha}, \mathbf{z})] \leq 0$, $\forall \boldsymbol{\alpha} \in \mathbb{A}_{\mathcal{F}}$. The reason for using Definition 2.1 is to ensure that the set of all feasible dual penalties $\mathbb{M}_{\mathcal{F}}$ is a function space (vector space). Note that Definition 2.1 does not exclude any “good” feasible dual penalties in Brown et al. (2010). If $M(\boldsymbol{\alpha}, \mathbf{z})$ is a feasible dual penalty in Brown et al. (2010), i.e., $\mathbb{E}_0 [M(\boldsymbol{\alpha}, \mathbf{z})] \leq 0$, $\forall \boldsymbol{\alpha} \in \mathbb{A}_{\mathcal{F}}$. Then $\mathbb{E}_0 [(M(\boldsymbol{\alpha}, \mathbf{z}) - \mathbb{E}_0 [M(\boldsymbol{\alpha}, \mathbf{z})])] = 0$, $\forall \boldsymbol{\alpha} \in \mathbb{A}_{\mathcal{F}}$. That is, $(M(\boldsymbol{\alpha}, \mathbf{z}) - \mathbb{E}_0 [M(\boldsymbol{\alpha}, \mathbf{z})])$ is a feasible dual penalty by Definition 2.1, and it always induces an upper bound as tight as the one induced by $M(\boldsymbol{\alpha}, \mathbf{z})$. In particular, Definition 2.1 does not exclude the optimal dual penalty (which will be defined in the following).

Brown et al. (2010) shows that the following dual representation of the primal DP (2.2), derived by subtracting a feasible dual penalty $M(\boldsymbol{\alpha}, \mathbf{z}) \in \mathbb{M}_{\mathcal{F}}$ from the objective function, relax the non-anticipativity constraint on all the feasible policies, and interchange the maximization and expectation, i.e.,

$$(D) : V_0^M(x_0) \triangleq \mathbb{E}_0 \left[\sup_{\mathbf{a} \in \mathcal{A}} \left\{ \sum_{n=0}^{N-1} r_n(x_n, a_n) + r_N(x_N) - M(\mathbf{a}, \mathbf{z}) \right\} \right], \quad (2.5)$$

yields an upper bound $V_0^M(x_0)$ on the optimal value function $V_0(x_0)$, where $\mathbf{a} := (a_0, \dots, a_{N-1})$ represents the action sequence and $\mathcal{A} := (\mathcal{A}_0, \dots, \mathcal{A}_{N-1})$. Conceptually, the dual problem (2.5) consists of a series of scenario-based deterministic pathwise optimization problems (referred to as inner optimization problems), in which the actions could be chosen with full knowledge of all outcomes of uncertainties. In practice, solving dual problem (2.5) is convenient via Monte Carlo simulation given that the inner optimization problems are tractable:

simulate multiple i.i.d. noise sequences \mathbf{z} , then solve the deterministic inner optimization problem corresponding to each \mathbf{z} , and finally take the average of the optimal values as an estimator of $V_0^M(x_0)$.

Note that a dual penalty is feasible if and only if it does not penalize any non-anticipative policy, we immediately have the weak duality. Furthermore, if minimizing $V_0^M(x_0)$ over all feasible dual penalties $M \in \mathbb{M}_{\mathcal{F}}$, then we have $\forall x_0 \in \mathcal{X}_0$,

$$\sup_{\alpha \in \mathbb{A}_{\mathcal{F}}} V_0^{\alpha}(x_0) = V_0(x_0) = \inf_{M \in \mathbb{M}_{\mathcal{F}}} V_0^M(x_0),$$

i.e., strong duality holds. In particular, the following proposition shows that the value-based optimal dual penalty in the form of (2.6) suffices for the strong duality to hold.

Proposition 2.1. (i) (**Weak Duality**) [Lemma 2.1 in Brown et al. (2010)] If $M(\alpha, \mathbf{z}) \in \mathbb{M}_{\mathcal{F}}$, then we have

$$V_0(x_0) \leq V_0^M(x_0).$$

(ii) (**Strong Duality**) [Theorem 2.3 in Brown et al. (2010)] Let $\{V_n(x_n)\}$ be the optimal value functions of the primal DP (2.2). Further let $M^*(\alpha, \mathbf{z})$ be the martingale difference sum of $\{V_n(x_n)\}$, i.e.,

$$M^*(\alpha, \mathbf{z}) \triangleq \sum_{n=0}^{N-1} (V_{n+1}(x_{n+1}) - \mathbb{E}[V_{n+1}(x_{n+1}) | x_n, a_n]), \quad \text{where } a_n = \alpha_n(x_n), \quad (2.6)$$

then $M^*(\alpha, \mathbf{z})$ is a feasible dual penalty. Moreover, it achieves strong duality, i.e.,

$$V_0^{M^*}(x_0) \triangleq \mathbb{E}_0 \left[\sup_{\alpha \in \mathbb{A}} \left\{ \sum_{n=0}^{N-1} r_n(x_n, \alpha_n(x_n)) + r_N(x_N) - M^*(\alpha, \mathbf{z}) \right\} \right] = V_0(x_0). \quad (2.7)$$

Note that the optimal dual penalty $M^*(\alpha, \mathbf{z})$ defined in (2.6) is a functional of (α, \mathbf{z}) because the state sequence (x_0, \dots, x_N) and action sequence \mathbf{a} depend on (α, \mathbf{z}) through the state dynamics (2.1) and the policy. Intuitively, weak duality suggests that the dual problem (2.5) with any feasible dual penalty could be used to generate an upper bound on the optimal value function, and strong duality implies that the advantage gained by the access to the future information is perfectly cancelled out in expectation by the optimal dual penalty. What is more striking about the optimal dual penalty is that the second equality in (2.7) is achieved almost surely for every inner optimization problem scenario. Therefore, we can drop the expectation sign in (2.7) and strengthen the result as

$$V_0(x_0) = \sup_{\alpha \in \mathbb{A}} \left\{ \sum_{n=0}^{N-1} r_n(x_n, \alpha_n(x_n)) + r_N(x_N) - M^*(\alpha, \mathbf{z}) \right\}, \quad a.s..$$

Therefore, in practical implementation, if the approximation of the optimal dual penalty is sufficiently accurate, then the variance of a one-sample estimator of the upper bound is small. Therefore, one only needs to solve a small number of deterministic inner optimization problems to generate a good upper bound estimator.

In theory, strong duality implies that the upper bound induced by the optimal dual penalty

is tight. However, one could hardly compute the optimal dual penalty exactly and close the duality gap since in general the optimal value functions $\{V_n(x_n)\}$ are not available. A naive alternative is to replace the optimal value functions with approximate ones $\{\tilde{V}_n(x_n)\}$ that might be induced by a suboptimal policy or certain approximate dynamic programming technique. The computational effort required in approximating the optimal dual penalty should be taken into consideration as well. A common method to estimate the conditional expectations in the optimal dual penalty (2.6) is nested simulation, which generates scenarios in the outer-layer and uses sample averaging in the inner-layer associated with each scenario. Hence, the total simulation effort is proportional to the number of outer-layer scenarios multiplied by the number of inner-layer samples. Albeit stable with a large-scale simulation, it might not be a desirable choice facing a limited computational budget. Several approaches have been developed in recent years that aim to generate good dual penalties without nested simulation. However, we note that two key things are missing in most of the existing approaches: the structure of the dual penalty space as well as the optimal dual penalty is not well-studied; the information such as the suboptimal policy and the sample paths generated in solving the primal problem is not well-utilized.

In the next section, we will present a framework of regression approach to approximating the optimal dual penalty that explores the structure of dual penalty space and efficiently reutilizes the information in the primal problem.

3 A Framework of Regression Approach

3.1 General Framework

The main idea of the regression approach is to view the optimal dual penalty as a point in the dual penalty space, then compute its coordinates w.r.t. a properly chosen functional basis, and finally estimate those coordinates in a non-nested manner. The key question lies in how to choose the functional basis and compute the corresponding coordinates of the optimal dual penalty. Let us first derive the general framework. Recall the optimal dual penalty

$$M^*(\boldsymbol{\alpha}, \mathbf{z}) = \sum_{n=0}^{N-1} (V_{n+1}(x_{n+1}) - \mathbb{E}[V_{n+1}(x_{n+1})|x_n, a_n]), \quad \text{where } a_n = \alpha_n(x_n).$$

Denote the n -th single martingale difference term in $M^*(\boldsymbol{\alpha}, \mathbf{z})$ by $M_n^*(\boldsymbol{\alpha}, \mathbf{z})$, i.e.,

$$M_n^*(\boldsymbol{\alpha}, \mathbf{z}) \triangleq (V_{n+1}(x_{n+1}) - \mathbb{E}[V_{n+1}(x_{n+1})|x_n, a_n]). \quad (3.1)$$

Then clearly $M^* = \sum_{n=0}^{N-1} M_n^*$. Further let $\mathcal{B} = \{b_i : i \in I\}$ denote a functional basis of $\mathbb{M}_{\mathcal{F}}$, where $\{b_i(\cdot)\}$ are functions on the support space Ξ of the noises $\{z_n\}$ such that $\int b_i(z)\rho(dz) = 0, \forall i \in I$, where recall that ρ is the probability measure induced by the noises $\{z_n\}$. Therefore, we have

$$\mathbb{E}[b_i(z_{n+1})] = 0, \quad \forall i \in I, n = 0, \dots, N-1. \quad (3.2)$$

Note that \mathcal{B} 's cardinality, $|I|$, might be countable or uncountable. We will discuss about the construction of \mathcal{B} or the properties it should possess later. Suppose we could express

$M_n^*(\boldsymbol{\alpha}, \mathbf{z})$ w.r.t. \mathcal{B} in the form of

$$(V_{n+1}(x_{n+1}) - \mathbb{E}[V_{n+1}(x_{n+1})|x_n, a_n]) = \sum_{i \in I} \beta_{n,i}(x_n, a_n) \cdot b_i(z_{n+1}), \quad (3.3)$$

where $\{\beta_{n,i}(x_n, a_n)\}$ are the coordinates. To compute $\{\beta_{n,i}(x_n, a_n)\}$, multiplying $b_i(z_{n+1})$ $\forall i \in I$ on both sides of (3.3) and taking conditional expectations w.r.t. \mathcal{F}_n , we obtain

$$\begin{aligned} & \mathbb{E}_n[V_{n+1}(x_{n+1}) \cdot b_i(z_{n+1})] - \mathbb{E}_n[\mathbb{E}[V_{n+1}(x_{n+1})|x_n, a_n] \cdot b_i(z_{n+1})] \\ &= \sum_{j \in I} \mathbb{E}_n[\beta_{n,j}(x_n, a_n) \cdot b_i(z_{n+1}) \cdot b_j(z_{n+1})]. \end{aligned} \quad (3.4)$$

Due to the Markov property, we have

$$\mathbb{E}_n[V_{n+1}(x_{n+1}) \cdot b_i(z_{n+1})] = \mathbb{E}[V_{n+1}(x_{n+1}) \cdot b_i(z_{n+1})|x_n, a_n].$$

Furthermore, by (3.2) and the facts that $\mathbb{E}[V_{n+1}(x_{n+1})|x_n, a_n]$ is \mathcal{F}_n -measurable and z_{n+1} is independent of \mathcal{F}_n , we have

$$\begin{aligned} & \mathbb{E}_n[\mathbb{E}[V_{n+1}(x_{n+1})|x_n, a_n] \cdot b_i(z_{n+1})] = \mathbb{E}[V_{n+1}(x_{n+1})|x_n, a_n] \cdot \mathbb{E}_n[b_i(z_{n+1})] \\ &= \mathbb{E}[V_{n+1}(x_{n+1})|x_n, a_n] \cdot \mathbb{E}[b_i(z_{n+1})] = 0. \end{aligned}$$

Hence, equation (3.4) could be rewritten as

$$\mathbb{E}[V_{n+1}(x_{n+1}) \cdot b_i(z_{n+1})|x_n, a_n] = \sum_{j \in I} \beta_{n,j}(x_n, a_n) \mathbb{E}[b_i(z_{n+1}) \cdot b_j(z_{n+1})], \quad (3.5)$$

where we use the facts that $\beta_{n,j}(x_n, a_n)$ is \mathcal{F}_n -measurable and z_{n+1} is independent of \mathcal{F}_n . Note that $\{\mathbb{E}[b_i(z_{n+1}) \cdot b_j(z_{n+1})]\}$ in (3.5) are constants that only depend on \mathcal{B} . Therefore, in principle we could view (3.5) as a system of linear equations with variables $\{\beta_{n,i}(x_n, a_n)\}$ (although the number of equations in the system, $|I|$, could be countably infinite or even uncountable). Assuming it is solvable (again this depends on the choice of \mathcal{B}), then we have

$$\beta_{n,i}(x_n, a_n) = \mathbb{E}[V_{n+1}(x_{n+1}) \cdot h_i(z_{n+1})|x_n, a_n], \quad \forall i \in I, n = 0, \dots, N-1, \quad (3.6)$$

where $\{h_i(\cdot) : i \in I\}$ are deterministic functions that only depend on \mathcal{B} because all the parameters of the linear system are uniquely determined by \mathcal{B} . It follows that the optimal dual penalty could be rewritten as

$$M^*(\boldsymbol{\alpha}, \mathbf{z}) = \sum_{n=0}^{N-1} \sum_{i \in I} \beta_{n,i}(x_n, a_n) \cdot b_i(z_{n+1}), \quad \text{where } a_n = \alpha_n(x_n), \quad (3.7)$$

where the coordinates $\{\beta_{n,i}(x_n, a_n)\}$ are given by (3.6).

Instead of approximating the optimal dual penalty $M^*(\boldsymbol{\alpha}, \mathbf{z})$ directly via nested simulation, now we are able to approximate it by estimating all its coordinates $\{\beta_{n,i}(x_n, a_n)\}$ w.r.t. \mathcal{B} . The difficulty lies in how to efficiently estimate the conditional expectations on the right hand side of (3.6). To avoid nested simulation, we propose the following regression scheme. The main idea is to treat $\beta_{n,i}(x_n, a_n)$ as the expected response, and

thus in the regression an observation of the expected response is a sample outcome of $V_{n+1}(x_{n+1}) \cdot h_i(z_{n+1})|x_n, a_n$, which is obtained by exercising a given suboptimal policy $\tilde{\alpha} := (\tilde{\alpha}_0, \tilde{\alpha}_1, \dots, \tilde{\alpha}_{N-1})$ and computing the value function along one sample path. Further let $\phi_{n,i}(x_n, a_n) \triangleq (\phi_{n,i}^1(x_n, a_n), \dots, \phi_{n,i}^K(x_n, a_n))^T$ denote the vector of regressors (dependent variables) in the regression, and later we will illustrate how to properly construct them based on the choice of \mathcal{B} . The linear regression model could be formulated as

$$V_{n+1}(x_{n+1}) \cdot h_i(z_{n+1})|x_n, a_n = (\phi_{n,i}(x_n, a_n))^T \theta_{n,i} + \epsilon_{n,i},$$

where $\theta_{n,i} := (\theta_{n,i}^1, \dots, \theta_{n,i}^K)^T$ is the vector of regression coefficients and $\epsilon_{n,i}$ is the noise. The complete algorithm is summarized as the following Algorithm 1.

Algorithm 1 A Regression-based Algorithm for Approximating the Optimal Dual Penalty

Input: Functional basis \mathcal{B} and suboptimal policy $\tilde{\alpha}$.

Output: $\tilde{M}(\alpha, \mathbf{z})$ —Approximation of the optimal dual penalty.

1. Initialization: Simulate M independent (state-action) sample paths under the given policy $\tilde{\alpha}$ subject to state dynamics (2.1). Denote the sample paths by

$$\{(x_0^j, a_0^j; x_1^j, a_1^j; \dots, x_{N-1}^j, a_{N-1}^j; x_N^j) : j = 1, \dots, M\},$$

where the noise sequences are $\{(z_1^j, z_2^j, \dots, z_N^j) : j = 1, \dots, M\}$ and $a_n^j = \tilde{\alpha}_n(x_n^j)$. Calculate the corresponding value functions along each sample path, denoted by

$$\{(V_0^j(x_0^j), V_1^j(x_1^j), \dots, V_N^j(x_N^j)) : j = 1, \dots, M\},$$

where $V_n^j(x_n^j) \triangleq \sum_{k=n}^{N-1} r_n(x_n^j, a_n^j) + r_N(x_N^j)$.

2. Iteration: For $n = 0, \dots, N-1$, $i \in I$, estimate the coordinate $\beta_{n,i}(x_n, a_n)$ by regression. The responses are $\{V_{n+1}^j(x_{n+1}^j) \cdot h_i(z_{n+1}^j) : j = 1, \dots, M\}$ and the design matrix is $(\phi_{n,i}^T(x_n^1, a_n^1), \dots, \phi_{n,i}^T(x_n^M, a_n^M))^T$, where $\phi_{n,i}(x_n^j, a_n^j) = (\phi_{n,i}^1(x_n^j, a_n^j), \dots, \phi_{n,i}^K(x_n^j, a_n^j))^T$. Let $\tilde{\theta}_{n,i} := (\tilde{\theta}_{n,i}^1, \dots, \tilde{\theta}_{n,i}^K)^T$ denote the vector of regression coefficients, i.e.,

$$\tilde{\theta}_{n,i} = \arg \min_{\theta_{n,i}} \frac{1}{M} \sum_{j=1}^M \left(V_{n+1}^j(x_{n+1}^j) \cdot h_i(z_{n+1}^j) - (\phi_{n,i}(x_n^j, a_n^j))^T \theta_{n,i} \right)^2.$$

3. Termination: Let

$$\tilde{\beta}_{n,i}(x_n, a_n) \triangleq (\phi_{n,i}(x_n, a_n))^T \tilde{\theta}_{n,i}.$$

be the estimation of $\beta_{n,i}(x_n, a_n)$ and

$$\tilde{M}(\alpha, \mathbf{z}) \triangleq \sum_{n=0}^{N-1} \sum_{i \in I} \tilde{\beta}_{n,i}(x_n, a_n) \cdot b_i(z_{n+1}), \quad \text{where } a_n = \alpha_n(x_n) \quad (3.8)$$

be the approximation of $M^*(\alpha, \mathbf{z})$.

For convenience, we refer to $\tilde{M}(\alpha, \mathbf{z})$ as the *regression-based (dual) penalty*. Note that one could implement Algorithm 1 efficiently by reusing the sample paths that were generated to

evaluate the suboptimal policy $\tilde{\alpha}$ in the primal problem. Hence, minimal extra simulation or computational costs are required. Moreover, generating a regression-based dual penalty $\tilde{M}(\alpha, \mathbf{z})$ as in (3.8) does not incur nested simulation, since the conditional expectations are estimated via regression instead. Hence, computational efficiency is significantly improved. Finally, one of the biggest advantages of the proposed regression approach over some of the existing approaches is that it preserves the dual feasibility of the optimal dual penalty as long as the function basis \mathcal{B} satisfies (3.2). That is, the regression-based penalty $\tilde{M}(\alpha, \mathbf{z})$ maintains to be a feasible dual penalty.

Theorem 3.1. *The regression-based dual penalty $\tilde{M}(\alpha, \mathbf{z})$ defined in (3.8) is a feasible dual penalty, i.e., $\mathbb{E}_0 [\tilde{M}(\alpha, \mathbf{z})] = 0, \forall \alpha \in \mathbb{A}_{\mathcal{F}}$.*

Proof. We have $\forall \alpha \in \mathbb{A}_{\mathcal{F}}$

$$\begin{aligned}
\mathbb{E}_0 [\tilde{M}(\alpha, \mathbf{z})] &= \mathbb{E} [\tilde{M}(\alpha, \mathbf{z}) | \mathcal{F}_0] = \mathbb{E} \left[\sum_{n=0}^{N-1} \sum_{i \in I} \tilde{\beta}_{n,i}(x_n, a_n) \cdot b_i(z_{n+1}) \middle| \mathcal{F}_0 \right] \\
&= \mathbb{E} \left[\sum_{n=0}^{N-2} \sum_{i \in I} \tilde{\beta}_{n,i}(x_n, a_n) \cdot b_i(z_{n+1}) \middle| \mathcal{F}_0 \right] + \mathbb{E} \left[\sum_{i \in I} \tilde{\beta}_{N-1,i}(x_{N-1}, a_{N-1}) \cdot b_i(z_N) \middle| \mathcal{F}_0 \right] \\
&\stackrel{(i)}{=} \mathbb{E} \left[\sum_{n=0}^{N-2} \sum_{i \in I} \tilde{\beta}_{n,i}(x_n, a_n) \cdot b_i(z_{n+1}) \middle| \mathcal{F}_0 \right] \\
&\quad + \mathbb{E} \left[\mathbb{E} \left[\sum_{i \in I} \tilde{\beta}_{N-1,i}(x_{N-1}, a_{N-1}) \cdot b_i(z_N) \middle| \mathcal{F}_{N-1} \right] \middle| \mathcal{F}_0 \right] \\
&\stackrel{(ii)}{=} \mathbb{E} \left[\sum_{n=0}^{N-2} \sum_{i \in I} \tilde{\beta}_{n,i}(x_n, a_n) \cdot b_i(z_{n+1}) \middle| \mathcal{F}_0 \right] + \mathbb{E} \left[\sum_{i \in I} \tilde{\beta}_{N-1,i}(x_{N-1}, a_{N-1}) \cdot \mathbb{E} [b_i(z_N) | \mathcal{F}_{N-1}] \middle| \mathcal{F}_0 \right] \\
&\stackrel{(iii)}{=} \mathbb{E} \left[\sum_{n=0}^{N-2} \sum_{i \in I} \tilde{\beta}_{n,i}(x_n, a_n) \cdot b_i(z_{n+1}) \middle| \mathcal{F}_0 \right] + \mathbb{E} \left[\sum_{i \in I} \tilde{\beta}_{N-1,i}(x_{N-1}, a_{N-1}) \cdot \mathbb{E} [b_i(z_N)] \middle| \mathcal{F}_0 \right] \\
&\stackrel{(iv)}{=} \mathbb{E} \left[\sum_{n=0}^{N-2} \sum_{i \in I} \tilde{\beta}_{n,i}(x_n, a_n) \cdot b_i(z_{n+1}) \middle| \mathcal{F}_0 \right] + 0 \\
&\dots \\
&= \mathbb{E} \left[\sum_{i \in I} \tilde{\beta}_{0,i}(x_0, a_0) \cdot b_i(z_1) \middle| \mathcal{F}_0 \right] = \sum_{i \in I} \tilde{\beta}_{0,i}(x_0, a_0) \cdot \mathbb{E} [b_i(z_1) | \mathcal{F}_0] \\
&= \sum_{i \in I} \tilde{\beta}_{0,i}(x_0, a_0) \cdot \mathbb{E} [b_i(z_1)] = 0,
\end{aligned}$$

where equality (i) follows from the tower property of conditional expectations and $\mathcal{F}_0 \subseteq \mathcal{F}_{N-1}$, equality (ii) follows from the fact that α is \mathcal{F} -adapted, and thus $\{\tilde{\beta}_{N-1,i}(x_{N-1}, a_{N-1})\}$ are \mathcal{F}_{N-1} -measurable, equality (iii) follows from the fact that z_N is independent of \mathcal{F}_{N-1} , and finally equality (iv) follows from (3.2). \square

From the proof, we can see the fact that each basis function $b_i(\cdot)$ has expectation zero w.r.t. the probability measure ρ is essential for $\tilde{M}(\alpha, \mathbf{z})$ to maintain the dual feasibility. Theorem

3.1 implies that the dual problem (2.5) with penalty $\widetilde{M}(\boldsymbol{\alpha}, \mathbf{z})$ provides a valid upper bound on the optimal value function, i.e., $V_0^{\widetilde{M}}(x_0) \geq V_0(x_0), \forall x_0 \in \mathcal{X}_0$. The tightness of the upper bound $V_0^{\widetilde{M}}(x_0)$ is directly affected by the accuracy of the penalty approximation, and it depends on the choice of \mathcal{B} as well as the accuracy of the regression. An ideal choice of \mathcal{B} should possess the following properties: completeness, orthogonality, and countability. The reason is that completeness guarantees (3.3) to hold, orthogonality significantly simplifies the linear system (3.5) and facilitates computation of the coordinates, and countability determines whether the linear system is solvable or not. Moreover, one would prefer a basis \mathcal{B} that satisfies (3.2) so that the resulted regression-based dual penalty is feasible, and thus induces valid upper bounds.

Following these guidelines, we consider two choices of \mathcal{B} as follows. The first one is the orthonormal basis of the Hilbert L^2 space induced by the probability measure ρ . We will show that it possesses good properties of completeness, orthogonality, and countability. The second one is the basis consisting of all the centralized moments of the noise distribution. We will show that it results in a simple linear system of equations for determining the coordinates. Finally, we will show that both functional bases result in regression-based penalties that are dual feasible and good approximations of the optimal dual penalty, and thus lead to valid and tight upper bounds numerically.

3.2 Regression Approach with L^2 Orthonormal Basis

To ease the presentation, let us first lay out some preliminaries for L^2 space. Recall the random noises $\{z_n\}$ are i.i.d. with probability measure ρ on support space $\Xi \subseteq \mathbb{R}^d$. Consider the L^2 space induced by the measure ρ , denoted by $L^2(\rho)$. In particular, $L^2(\rho)$ is the space consisting of all ρ -measurable scalar function f such that $|f|^2$ is integrable. That is, $L^2(\rho) = \{f : \int_{\Xi} |f|^2 d\rho < \infty\}$.

By Theorem 4.1.3 in Bogachev (2007), $L^2(\rho)$ is a Banach (hence complete) space equipped with the L^2 -norm $\|\cdot\|_2$, where

$$\|f\|_2 = \sqrt{\int_{\Xi} |f|^2 d\rho}, \quad \forall f \in L^2(\rho).$$

Here note that a metric space S is called complete if every Cauchy sequence of points in S has a limit in S . A nice property of $L^2(\rho)$ is that the L^2 norm $\|\cdot\|_2$ is generated by the inner product defined by

$$\langle f, g \rangle \triangleq \int_{\Xi} f \cdot g d\rho, \quad \forall f, g \in L^2(\rho).$$

Obviously, $\|f\|_2 = \sqrt{\langle f, f \rangle}$. Therefore, $L^2(\rho)$ is a Hilbert space (see, e.g., page 255 in Bogachev (2007)) equipped with the inner product defined above. Now let us introduce the definition of *orthonormal basis* in a Hilbert space.

Definition 3.1. (Orthonormal Basis) For a Hilbert space H equipped with inner product $\langle \cdot, \cdot \rangle_H$ and norm $\|\cdot\|_H$, a family of mutually orthogonal unit vectors $\mathcal{E} = \{e_i \in H : i \in I\}$ is called an (complete) orthonormal basis of H if

(i) $\|e_i\|_H = 1, \forall i \in I; \langle e_i, e_j \rangle_H = 0, \forall i, j \in I, i \neq j.$

(ii) Every element $f \in H$ can be represented as the linear combination of $\{e_i\}$, i.e.,

$$f \stackrel{H}{=} \sum_{i \in I} \beta_i e_i, \text{ or equivalently, } \left\| f - \sum_{i \in I} \beta_i e_i \right\|_H = 0,$$

where at most countably many coefficients β_i may be nonzero.

It immediately follows that if $f \stackrel{H}{=} \sum_i \beta_i e_i$, then

$$\beta_i = \langle f, e_i \rangle_H \quad \text{and} \quad \|f\|_H = \sqrt{\sum_i |\beta_i|^2}.$$

For a general Hilbert space H , by Zorns lemma (see, e.g., Bogachev (2007)), there always exists an orthonormal basis; however, the number of basis functions in that orthonormal basis might be uncountable. Fortunately, for the specific Hilbert space $L^2(\rho)$, we have the following proposition on the countability of its orthonormal basis.

Proposition 3.1. *Let $\mathcal{E} = \{e_i \in L^2(\rho) : i \in I\}$ be an orthonormal basis of $L^2(\rho)$, then \mathcal{E} is countable.*

Proof. Note that a metric space S is called separable if it contains a countable and everywhere dense subset. By Corollary 4.2.2 in Bogachev (2007), $L^2(\rho)$ is a separable Hilbert space. Therefore, by Corollary 4.3.4 in Bogachev (2007), any orthonormal basis of $L^2(\rho)$ is countable. \square

Proposition 3.1 implies that we can rewrite the orthonormal basis \mathcal{E} as $\mathcal{E} = \{e_i \in L^2(\rho) : i = 0, 1, 2, \dots\}$. Without loss of generality, let us assume $e_0 \equiv 1$; otherwise we could simply perform the Gram–Schmidt process to achieve that. Note that the basis \mathcal{E} is uniquely determined when $e_0 \equiv 1$. To gain more intuition, below are examples of orthonormal basis of $L^2(\rho)$ for several common probability measures.

Example 1. Finite Discrete Distribution. Assume ρ is a finite discrete probability measure with positive probability masses on $\Xi = \{y_0, \dots, y_{p-1}\}$. It is easy to verify that the cardinality of any orthonormal basis of $L^2(\rho)$ is p . In fact,

$$\mathcal{G} = \{g_i : g_i(y) = \frac{1}{\sqrt{\rho(y_i)}} \mathbb{1}\{y = y_i\}, i = 0, 1, \dots, p-1.\}$$

is an orthonormal basis of $L^2(\rho)$.

Example 2. Standard Normal Distribution. Assume ρ is the measure induced by a standard normal distribution such that $\rho(z) = 1/\sqrt{2\pi} \exp(-\frac{z^2}{2}), \forall z \in \Xi = \mathbb{R}$, then the Hermite polynomials

$$\{1; z; \frac{1}{\sqrt{2}}(z^2 - 1); \frac{1}{\sqrt{6}}(z^3 - 3z); \dots\}$$

is a countable orthonormal basis of $L^2(\rho)$.

Next, we will show that the L^2 orthonormal basis $\mathcal{E}_1 := \mathcal{E} \setminus \{e_0\}$ is a valid choice of basis \mathcal{B} in the proposed regression approach. In view of the state dynamics (2.1), for a fixed (x_n, a_n) pair, $V_{n+1}(x_{n+1}) = V_{n+1}(f(x_n, a_n, z_{n+1}))$ is a random variable as a function of the random variable z_{n+1} . Further suppose $V_{n+1}(x_{n+1})$ has finite second moment, i.e., $V_{n+1}(f(x_n, a_n, \cdot)) \in L^2(\rho)$. Since $\mathcal{E} = \{e_i : i = 0, 1, \dots\}$ with $e_0 \equiv 1$ is an orthonormal basis of $L^2(\rho)$, we can express $V_{n+1}(x_{n+1})$ w.r.t. \mathcal{E} , i.e.,

$$V_{n+1}(x_{n+1}) = \sum_{i=0}^{\infty} \delta_{n,i}(x_n, a_n) \cdot e_i(z_{n+1}), \quad (3.9)$$

where, following Definition 3.1, $\{\delta_{n,i}(x_n, a_n)\}$ are coordinates such that

$$\begin{aligned} \delta_{n,i}(x_n, a_n) &= \langle V_{n+1}(x_{n+1}), e_i(z_{n+1}) \rangle = \int_{\Xi} V_{n+1}(x_{n+1}) \cdot e_i(z_{n+1}) \rho(dz_{n+1}) \\ &= \mathbb{E}[V_{n+1}(x_{n+1}) \cdot e_i(z_{n+1}) | x_n, a_n], \quad i = 0, 1, \dots \end{aligned} \quad (3.10)$$

In particular, noting that $e_0 \equiv 1$, then

$$\delta_{n,0}(x_n, a_n) = \mathbb{E}[V_{n+1}(x_{n+1}) | x_n, a_n]. \quad (3.11)$$

Combining (3.1), (3.9) and (3.11), we have

$$\begin{aligned} M_n^*(\boldsymbol{\alpha}, \mathbf{z}) &= (V_{n+1}(x_{n+1}) - \mathbb{E}[V_{n+1}(x_{n+1}) | x_n, a_n]) \\ &= \sum_{i=0}^{\infty} \delta_{n,i}(x_n, a_n) \cdot e_i(z_{n+1}) - \delta_{n,0}(x_n, a_n) \cdot e_0(z_{n+1}) \\ &= \sum_{i=1}^{\infty} \delta_{n,i}(x_n, a_n) \cdot e_i(z_{n+1}). \end{aligned}$$

Thus, the optimal dual penalty could be rewritten as

$$M^*(\boldsymbol{\alpha}, \mathbf{z}) = \sum_{n=0}^{N-1} \sum_{i=1}^{\infty} \delta_{n,i}(x_n, a_n) \cdot e_i(z_{n+1}), \quad \text{where } a_n = \alpha_n(x_n), \quad (3.12)$$

where $\delta_{n,i}(x_n, a_n)$ is given by (3.10). Comparing (3.7) with (3.12), we can see that the L^2 orthonormal basis $\mathcal{E}_1 = \{e_1, e_2, \dots\}$ is a valid choice of \mathcal{B} in the regression approach. Moreover, the function $h_i(\cdot)$ in (3.6) reduces to $e_i(\cdot)$ due to \mathcal{E}_1 's orthogonality. Hence, the corresponding coordinates $\{\delta_{n,i}(x_n, a_n)\}$ are straightforward to compute as in (3.10).

Finally, let us show that using \mathcal{E}_1 in the regression approach induces feasible dual penalties, and thus valid upper bounds on the optimal value function. To be more specific, let

$$\widetilde{M}^l(\boldsymbol{\alpha}, \mathbf{z}) \triangleq \sum_{n=0}^{N-1} \sum_{i=1}^P \widetilde{\delta}_{n,i}(x_n, a_n) \cdot e_i(z_{n+1}), \quad \text{where } a_n = \alpha_n(x_n),$$

denote the regression-based dual penalty w.r.t. the basis \mathcal{E}_1 , where P is the order of the basis truncation, and $\{\widetilde{\delta}_{n,i}(x_n, a_n)\}$ is the estimation of $\{\delta_{n,i}(x_n, a_n)\}$ after regression. Notice that $\forall n = 0, \dots, N-1, i = 1, 2, \dots$

$$\mathbb{E}[e_i(z_{n+1})] = \int_{\Xi} e_i(z) \rho(dz) = \int_{\Xi} e_i(z) \cdot e_0(z) \rho(dz) = \langle e_i(z_{n+1}), e_0(z_{n+1}) \rangle = 0, \quad (3.13)$$

where we use the fact $e_0(z_{n+1}) \equiv 1$ and \mathcal{E} 's orthogonality. Then \mathcal{E}_1 is a choice of \mathcal{B} that satisfies (3.2). By Theorem 3.1, $\widetilde{M}^l(\boldsymbol{\alpha}, \mathbf{z})$ is dual feasible.

As mentioned previously, the tightness of the upper bound $V_0^{\widetilde{M}^l}(x_0)$ is directly affected the accuracy of the approximation $\widetilde{M}^l(\boldsymbol{\alpha}, \mathbf{z})$. If zero regression error is assumed, then $\widetilde{M}^l(\boldsymbol{\alpha}, \mathbf{z})$ converges to $M^*(\boldsymbol{\alpha}, \mathbf{z})$ in L^2 as the order of basis truncation, P , goes to infinity, due to the completeness of \mathcal{E} . When regression error is taken into account, the construction of regressors ϕ is very important in controlling that error. We will discuss this issue later.

Although, in theory, the unique L^2 orthonormal basis \mathcal{E}_1 could be calculated sequentially via Gram-Schmidt process for any probability measure ρ . In practice, this procedure might be complicated for complex probability measures, which is the drawback of using L^2 orthonormal basis in the regression approach. Next, we will derive an alternative functional basis that is easy to calculate. It is derived by carrying out Taylor series expansion on the value function, and consists of centralized moments of the noise distribution along each dimension.

3.3 Regression Approach with Taylor Series Basis

Now let us derive the functional basis of $\mathbb{M}_{\mathcal{F}}$ induced by Taylor series expansion on the value function. To ease the presentation, let us further assume the state dynamics is linear in noise as follows.

$$x_{n+1} = b(x_n, a_n) + \sigma(x_n, a_n)z_{n+1}, \quad n = 0, 1, \dots, N-1, \quad (3.14)$$

where $b(\cdot, \cdot)$ and $\sigma(\cdot, \cdot)$ are deterministic \mathcal{F} -measurable functions, and z_{n+1} has finite moments (up to some order). Although the derivation could be easily extended to the setting with general state dynamics (2.1), we confine it within (3.14) to better demonstrate that several existing approaches could be regarded as special cases of the proposed approach. Note that the state dynamics (3.14) covers a wide range of dynamic programming problems. For example, it has been widely used when modeling dynamics of asset/stock prices in finance literature, where, for instance, x might represent the asset/stock prices, $b(\cdot, \cdot)$ and $\sigma(\cdot, \cdot)$ might represent the drifts and the volatilities in the asset prices, respectively. The reward functions and the objective remain the same. Therefore, the primal problem is solved as before. We will focus on deriving a functional basis of the dual penalty space that facilitates approximation of the optimal dual penalty.

For convenience, let $\widehat{x}_{n+1} := b(x_n, a_n)$, i.e., \widehat{x}_{n+1} represents the expected future state at period $(n+1)$ given state x_n and action a_n at period n . In particular, \widehat{x}_{n+1} is \mathcal{F}_n -measurable. Rewrite the state dynamics (3.14) as

$$x_{n+1} - \widehat{x}_{n+1} = \sigma(x_n, a_n)z_{n+1}, \quad n = 0, 1, \dots, N-1.$$

Consider the Taylor series expansion on $V_{n+1}(x_{n+1})$ around the point \widehat{x}_{n+1} up to R -th order

$$V_{n+1}(x_{n+1}) \approx V_{n+1}(\widehat{x}_{n+1}) + \sum_{r=1}^R \frac{1}{r!} \frac{\partial^r}{\partial x^r} V_{n+1}(\widehat{x}_{n+1}) \cdot (\sigma)^r (z_{n+1})^r, \quad (3.15)$$

where we use σ as an abbreviation for $\sigma(x_n, a_n)$. Taking expectations w.r.t. \mathcal{F}_n on both sides of (3.15), we have

$$\mathbb{E}[V_{n+1}(x_{n+1})|x_n, a_n] \approx V_{n+1}(\hat{x}_{n+1}) + \sum_{r=1}^R \frac{1}{r!} \frac{\partial^r}{\partial x^r} V_{n+1}(\hat{x}_{n+1}) \cdot (\sigma)^r \mathbb{E}[(z_{n+1})^r], \quad (3.16)$$

where we use the fact that $V_{n+1}(\hat{x}_{n+1})$, its partial derivatives and $\sigma(x_n, a_n)$ are \mathcal{F}_n -measurable, and z_{n+1} is independent of \mathcal{F}_n . Subtracting (3.16) from (3.15), we have

$$V_{n+1}(x_{n+1}) - \mathbb{E}[V_{n+1}(x_{n+1})|x_n, a_n] \approx \sum_{r=1}^R \frac{1}{r!} \frac{\partial^r}{\partial x^r} V_{n+1}(\hat{x}_{n+1}) \cdot (\sigma)^r \left((z_{n+1})^r - \mathbb{E}[(z_{n+1})^r] \right).$$

The above approximation has a nice structure because the expectations are on $(z_{n+1})^r$ instead of $V_{n+1}(x_{n+1})$, which could be calculated analytically. However, the partial derivatives could be difficult to compute directly and often require approximation methods such as finite difference. Nevertheless, it inspires us to consider the functional basis

$$\mathcal{D} \triangleq \{d_r : d_r(z) = (z^r - \mathbb{E}[z^r]), r = 1, \dots, R\}$$

of the dual penalty space. It follows that the optimal dual penalty $M^*(\boldsymbol{\alpha}, \mathbf{z})$ could be expressed w.r.t. \mathcal{D} as

$$M^*(\boldsymbol{\alpha}, \mathbf{z}) \approx \sum_{n=0}^{N-1} \sum_{r=1}^R \gamma_{n,r}(x_n, a_n) \cdot \left((z_{n+1})^r - \mathbb{E}[(z_{n+1})^r] \right) = \sum_{n=0}^{N-1} \sum_{r=1}^R \gamma_{n,r}(x_n, a_n) \cdot d_r(z_{n+1}),$$

where $\{\gamma_{n,r}(x_n, a_n)\}$ are the corresponding coordinates. Following the derivations in (3.3)-(3.5), we obtain a simple linear system of equations for $\{\gamma_{n,r}(x_n, a_n)\}$ as follows:

$$\begin{bmatrix} \mathbb{E}_n[V_{n+1}(x_{n+1})(z_{n+1} - \mathbb{E}[z_{n+1}])] \\ \mathbb{E}_n[V_{n+1}(x_{n+1})((z_{n+1})^2 - \mathbb{E}[(z_{n+1})^2])] \\ \vdots \\ \mathbb{E}_n[V_{n+1}(x_{n+1})((z_{n+1})^R - \mathbb{E}[(z_{n+1})^R])] \end{bmatrix} \approx \begin{bmatrix} Cov \begin{pmatrix} z_{n+1} \\ (z_{n+1})^2 \\ \vdots \\ (z_{n+1})^R \end{pmatrix} \end{bmatrix} \begin{bmatrix} \gamma_{n,1}(x_n, a_n) \\ \gamma_{n,2}(x_n, a_n) \\ \vdots \\ \gamma_{n,R}(x_n, a_n) \end{bmatrix}, \quad (3.17)$$

where $Cov(\cdot)$ represents the covariance matrix of a random vector. Note that it is nonsingular, and thus we can solve (3.17) as

$$\begin{bmatrix} \gamma_{n,1}(x_n, a_n) \\ \gamma_{n,2}(x_n, a_n) \\ \vdots \\ \gamma_{n,R}(x_n, a_n) \end{bmatrix} \approx \begin{bmatrix} Cov \begin{pmatrix} z_{n+1} \\ (z_{n+1})^2 \\ \vdots \\ (z_{n+1})^R \end{pmatrix} \end{bmatrix}^{-1} \begin{bmatrix} \mathbb{E}_n[V_{n+1}(x_{n+1})(z_{n+1} - \mathbb{E}[z_{n+1}])] \\ \mathbb{E}_n[V_{n+1}(x_{n+1})((z_{n+1})^2 - \mathbb{E}[(z_{n+1})^2])] \\ \vdots \\ \mathbb{E}_n[V_{n+1}(x_{n+1})((z_{n+1})^R - \mathbb{E}[(z_{n+1})^R])] \end{bmatrix}. \quad (3.18)$$

For simplicity, let us rewrite (3.18) explicitly as

$$\gamma_{n,r}(x_n, a_n) \approx \mathbb{E}[V_{n+1}(x_{n+1})l_r(z_{n+1})|x_n, a_n], \quad r = 1, \dots, R, \quad (3.19)$$

where $\{l_r(\cdot)\}$ are simple polynomial functions of degrees $\leq R$ because the vector on the right hand side of (3.18) only involves monomials of z_{n+1} of degrees $\leq R$. Therefore, the

coordinates $\{\gamma_{n,r}(x_n, a_n)\}$ could be easily computed, which is the main advantage of using the function basis \mathcal{D} in the regression approach. The drawback is that it does not guarantee convergence of the optimal dual penalty approximation as R goes to infinity, because in general Taylor series expansion is not convergent.

Finally, similar to using the function basis \mathcal{E}_1 , let us show that using \mathcal{D} in the regression approach produces feasible dual penalties, and thus valid upper bounds on the optimal value function. To this end, let

$$\widetilde{M}^t(\boldsymbol{\alpha}, \mathbf{z}) \triangleq \sum_{n=0}^{N-1} \sum_{r=1}^R \widetilde{\gamma}_{n,i}(x_n, a_n) \cdot d_i(z_{n+1}), \quad \text{where } a_n = \alpha_n(x_n)$$

denote the regression-based dual penalty generated via Algorithm 1 with \mathcal{D} as the functional basis \mathcal{B} , where $\{\widetilde{\gamma}_{n,i}(x_n, a_n)\}$ is the estimation of $\{\gamma_{n,i}(x_n, a_n)\}$ through regression. Note that for $n = 0, \dots, N-1, r = 1, \dots, R$,

$$\mathbb{E}[d_r(z_{n+1})] = \mathbb{E}[(z_{n+1})^R - \mathbb{E}[(z_{n+1})^R]] = \mathbb{E}[(z_{n+1})^R] - \mathbb{E}[(z_{n+1})^R] = 0.$$

Thus, \mathcal{D} is a choice of \mathcal{B} that satisfies (3.2). Therefore, by Theorem 3.1 $\widetilde{M}^t(\boldsymbol{\alpha}, \mathbf{z})$ is a feasible dual penalty. Similarly, an immediate implication is that the dual problem (2.5) with penalty $\widetilde{M}^t(\boldsymbol{\alpha}, \mathbf{z})$ provides a valid upper bound on the optimal value function, i.e., $V_0^{\widetilde{M}^t}(x_0) \geq V_0(x_0), \forall x_0 \in \mathcal{X}_0$. Again, the tightness of the upper bound $V_0^{\widetilde{M}^t}(x_0)$ is directly affected by the accuracy of the approximation $\widetilde{M}^t(\boldsymbol{\alpha}, \mathbf{z})$, which heavily relies on the regressors $\boldsymbol{\phi}$ constructed in the regression. We will discuss this issue next.

Remark 3.1. *Interestingly, one could verify that if the noises $\{z_n\}$ are symmetric, then the first two basis functions $d_1, d_2 \in \mathcal{D}$ coincide with the first two basis functions $e_1, e_2 \in \mathcal{E}_1$ (after taking out the normalization factors), respectively. For example, when $\{z_n\}$ follows a standard normal distribution, then $\{d_1 = z, d_2 = (z^2 - 1)\} \subset \mathcal{D}$ coincides with $\{e_1 = z, e_2 = \frac{1}{\sqrt{2}}(z^2 - 1)\} \subset \mathcal{E}_1$ in Example 2 of Section 3.2.*

3.4 Implementation

We will discuss two important implementation issues regarding the proposed regression approach, i.e., how to properly construct the regressors in the regression and how to solve the inner optimization problem in the resulted dual problem.

As mentioned previously, the accuracy of the regression-based dual penalty $\widetilde{M}(\boldsymbol{\alpha}, \mathbf{z})$ heavily relies on the choice of regressors $\boldsymbol{\phi}$ for regressing the coordinates β . The general guideline is to construct $\boldsymbol{\phi}$ similar to β in structure. In view of the equation $\beta_{n,i}(x_n, a_n) = \mathbb{E}[V_{n+1}(x_{n+1}) \cdot h_i(z_{n+1}) | x_n, a_n]$, a natural candidate regressor is $\mathbb{E}[\widetilde{V}_{n+1}(x_{n+1}) \cdot h_i(z_{n+1}) | x_n, a_n]$, where $\widetilde{V}_{n+1}(\cdot)$ is an approximation of $V_{n+1}(\cdot)$ with a closed-form expression. For example, $\widetilde{V}_{n+1}(\cdot)$ could be the value function induced by a naive policy, or an approximate optimal value function, or the value function of a simpler DP (e.g., through a relaxation of the constraints on the actions of the original DP). Multiple regressors of such form could be included in the regression. More specifically, if the Taylor series basis \mathcal{D} is used, then there exists good regressors of a simpler form. Note that the coordinates $\{\gamma_{n,r}(x_n, a_n)\}$ take the positions

of the partial derivatives in the Taylor series expansion of the optimal dual penalty. Hence, a good candidate regressor is $\frac{\partial^r}{\partial x^r} \tilde{V}_{n+1}(\hat{x}_{n+1}) \cdot (\sigma)^r$, where $\tilde{V}_{n+1}(\cdot)$ admits partial derivatives that are easy to derive. Again, $\tilde{V}_{n+1}(\cdot)$ could be the value function induced by a naive policy, or an approximate optimal value function, or the value function of a simpler DP.

Theoretically, after plugging the regression-based dual penalty $\tilde{M}(\boldsymbol{\alpha}, \mathbf{z})$ into the dual problem (2.5), we only need to solve a series of deterministic inner optimization problem to generate an upper bound estimator. However, they might be hard to solve without good structural properties such as convexity. Therefore, for a convex DP, ideally the dual penalty plugged in the dual problem should preserve convexity in the inner optimization problem. Note that the optimal dual penalty $M^*(\boldsymbol{\alpha}, \mathbf{z})$ might not be convex even if the value functions $\{V_n(x_n)\}$ are concave, and in this case plugging in the optimal dual penalty causes the inner optimization problem to lose convexity. One solution to this issue is to linearize the optimal dual penalty around a fixed policy. The resulted optimal dual penalty approximation is affine in policy, and thus preserves convexity in the dual problem.

The proposed framework of regression approach does not suffer from this issue. It is robust in the sense that an optimal dual penalty approximation with desired structural properties could be generated by using proper regressors, noting that the resulted dual penalty is a linear combination of the regressors. For example, for a convex DP, the regression-based dual penalty preserves convexity in the dual problem by constructing regressors that are affine in policy.

4 Special Cases

It turns out, in a broader sense, several existing approaches to approximating the optimal dual penalty are special cases of the proposed regression approach under specific settings.

4.1 American Option Pricing

Optimal stopping problem is one of the simplest non-trivial types of DPs. Various methodologies and numerical methods are first proposed under the setting of optimal stopping, and then generalized to general DPs. American option pricing, as one of the most salient applications of optimal stopping, has been extensively studied. We will demonstrate that the non-nested simulation approach developed in Belomestny et al. (2009) to approximating the optimal dual martingale in American option pricing could be regarded as a special case of the proposed regression approach with the first-order Taylor series basis $\{d_1\}$.

The problem setting is as follows. Assume the asset price $X(t)$ satisfies a stochastic differential equation (SDE) w.r.t. Brownian motion

$$dX(t) = b(t, X(t))dt + \sigma(t, X(t))dW(t),$$

where $t \in [0, T]$, $X(t)$ denotes the asset price at time t with given initial deterministic value $X(0) = X_0$, $W(t)$ represents the (standard) Brownian motion, and the coefficients b, σ are functionals satisfying mild regularity conditions. By convention, we use $\{\mathcal{F}_t : 0 \leq t \leq T\}$ to denote the augmented information filtration generated by the Brownian motion $\{W(t)\}$.

Consider a Bermudan option on $X(t)$ that can be exercised at any date from the time set $\mathcal{T} = \{T_0, T_1, \dots, T_N\}$, with $0 = T_0 < T_1 < \dots < T_N = T$. When exercising at time $T_n \in \mathcal{T}$, the option holder receives a payoff $H_{T_n} = h(T_n, X(T_n))$. The goal is to evaluate the price of the Bermudan option, that is, to find

$$(P) : V_0 = \sup_{\tau \in \mathcal{T}} \mathbb{E} [h(\tau, X(\tau)) | X(0) = X_0],$$

where τ is an exercising strategy, i.e., a stopping time adapted to the filtration $\mathcal{F} = \{\mathcal{F}_{T_n} : n = 0, \dots, N\}$ taking values in \mathcal{T} , and V_0 denotes the Bermudan option price at time T_0 with initial asset price X_0 .

As for the dual formulation, Anderson and Broadie (2004) and Haugh and Kogan (2004) show that all the \mathcal{F} -adapted martingales in $\mathbb{M}_{\mathcal{F}} = \{M = (M_0, \dots, M_N) : M \text{ is } \mathcal{F}\text{-adapted}\}$ are feasible dual penalties. In particular, the optimal dual martingale M^* is the Doob-Meyer martingale component of the Bermudan price process $\{V_{T_n} : n = 0, \dots, N\}$, i.e.,

$$\begin{cases} M_0^* = 0, \\ M_{T_{n+1}}^* = M_{T_n}^* + V_{T_{n+1}} - \mathbb{E}_{T_n} [V_{T_{n+1}}], n = 0, \dots, N-1. \end{cases} \quad (4.1)$$

The question is how to efficiently estimate M^* in a non-nested manner. Note that the martingale M^* is driven by Brownian motion, by martingale representation theorem we have

$$M_{T_n}^* = \int_0^{T_n} C_s dW(s), \quad n = 0, \dots, N,$$

where C_s is a predictable process w.r.t. the filtration \mathcal{F} . Belomestny et al. (2009) proposes an Ito sum approximation scheme as follows.

$$M_{T_n}^* = \int_0^{T_n} C_s dW(s) \approx \sum_{i=0}^{n-1} C_{T_i} (\Delta W_{T_{i+1}}), \quad n = 0, \dots, N,$$

where $(\Delta W_{T_{i+1}}) := (W(T_{i+1}) - W(T_i))$ denotes the Brownian increment. Combining with (4.1), we have

$$V_{T_{n+1}} - \mathbb{E}_{T_n} [V_{T_{n+1}}] = M_{T_{n+1}}^* - M_{T_n}^* \approx C_{T_n} (\Delta W_{T_{n+1}}), \quad n = 0, \dots, N-1. \quad (4.2)$$

Multiplying both sides of (4.2) by the Brownian increment $(\Delta W_{T_{n+1}})$ and taking conditional expectations w.r.t. \mathcal{F}_{T_n} , we obtain

$$C_{T_n} \approx \frac{1}{T_{n+1} - T_n} \mathbb{E}_{T_n} [V_{T_{n+1}} (\Delta W_{T_{n+1}})], \quad n = 0, \dots, N-1.$$

Starting from here, regression is further applied to estimate C_{T_n} in order to avoid nested simulation. Eventually, a regression-based dual martingale $\widetilde{M} = (\widetilde{M}_{T_0}, \widetilde{M}_{T_1}, \dots, \widetilde{M}_{T_N})$ is given by

$$\widetilde{M}_{T_n} \triangleq \sum_{i=0}^{n-1} \widetilde{C}_{T_i} \cdot (\Delta W_{T_{i+1}}), \quad n = 0, \dots, N, \quad (4.3)$$

where \widetilde{C}_{T_i} is the estimation of C_{T_i} after regression. Since the celebrated Ito's Lemma and Martingale Representation Theorem could be viewed as the results of carrying out Taylor

series expansion on a stochastic function, an Ito sum approximation of a stochastic integral could be viewed as a first-order Taylor series expansion type of scheme. Thus, approximation scheme (4.3) could be viewed as a special case of the proposed regression approach with the first-order Taylor series basis $\{d_1\}$.

Remark 4.1. *Belomestny et al. (2009) actually considers a finer Ito sum approximation of the optimal dual martingale by introducing a finer partition of the time span. Here we present a simpler version of their approach so that its connection with our proposed approach is more clear.*

4.2 Controlled Markov Diffusion

Ye and Zhou (2015) studies the form of optimal dual penalty under the setting of controlled Markov diffusion (CMD) and show that it is a stochastic integral. It then inspires the authors to propose an approximation scheme for the optimal dual penalty of a discrete-time DP with state dynamics (3.14) as follows:

$$\widehat{M}(\boldsymbol{\alpha}, \mathbf{z}) \triangleq \sum_{n=0}^{N-1} \frac{\partial}{\partial x} V_n(x_n) \cdot \sigma(x_n, a_n) \cdot z_{n+1}, \quad (4.4)$$

where the partial derivative $\frac{\partial}{\partial x} V_n(x_n)$ is estimated via finite difference method. This scheme is derived by mimicking the Ito sum approximation of the optimal dual penalty (which is a stochastic integral) under the setting of CMD. The connection with our proposed regression approach is evident in the sense that the approximation scheme (4.4) could be interpreted as a result of a first-order Taylor series expansion of the optimal dual penalty, and the coordinates are estimated using finite difference method instead of regression. Another difference we must point out is that $\frac{\partial}{\partial x} V_{n+1}(\hat{x}_{n+1})$ instead of $\frac{\partial}{\partial x} V_n(x_n)$ is used in our proposed regression approach. We will show that, at least for the following linear quadratic control (LQC) problem, our approach is better in terms of approximation accuracy.

4.3 Linear Quadratic Control

Now let us consider the Taylor series expansion of the optimal dual penalty for the classic LQC problem, since the LQC problem has been extensively studied and applied, see Bertsekas (2007). Assuming a linear state dynamics

$$x_{n+1} = A_n x_n + B_n a_n + z_{n+1}, \quad n = 0, \dots, N-1,$$

with an expected quadratic cost

$$\mathbb{E}_0 \left[x_N^T Q_N x_N + \sum_{n=0}^{N-1} (x_n^T Q_n x_n + a_n^T R_n a_n) \right], \quad (4.5)$$

where A_n, B_n, Q_n, R_n are known matrices of appropriate dimensions, $\{z_n\}$ are proper-dimensional i.i.d. vector random variables with zero mean and finite second moment, and Q_n and R_n are positive semi-definite. The objective is to select a non-anticipative policy $\boldsymbol{\alpha}^*$ that minimizes the total costs defined in (4.5). It turns out that the optimal control policy (e.g., see

Bertsekas (2007)) α_n^* admits a closed-form expression $\alpha_n^*(x_n) = L_n x_n$, where the gain matrix L_n is given by $L_n = -(B_n^T K_{n+1} B_n + R_n)^{-1} B_n^T K_{n+1} A_n$, and the positive semi-definite matrix K_n is given recursively by

$$\begin{cases} K_N = Q_N, \\ K_n = A_n^T \left(K_{n+1} - K_{n+1} B_n (B_n^T K_{n+1} B_n + R_n)^{-1} B_n^T K_{n+1} \right) A_n + Q_n. \end{cases}$$

The optimal value function $V_n(x_n)$ then takes the form

$$V_n(x_n) = x_n^T K_n x_n + \sum_{i=n}^{N-1} \mathbb{E} [z_{n+1}^T K_{n+1} z_{n+1}], \quad n = 0, \dots, N.$$

For the corresponding dual problem via information relaxation, it is easy to verify that (see, e.g., Haugh and Lim (2012)) the value-based optimal dual penalty $M^*(\alpha, \mathbf{z})$ takes the form

$$\begin{aligned} M^*(\alpha, \mathbf{z}) &= \sum_{n=0}^{N-1} \left(V_{n+1}(x_{n+1}) - \mathbb{E} [V_{n+1}(x_{n+1}) | x_n, a_n] \right) \\ &= \sum_{n=0}^{N-1} \left\{ 2(A_n x_n + B_n a_n)^T K_{n+1} z_{n+1} + z_{n+1}^T K_{n+1} z_{n+1} - \mathbb{E} [z_{n+1}^T K_{n+1} z_{n+1}] \right\}. \end{aligned}$$

Now let us approximate $M^*(\alpha, \mathbf{z})$ using the second-order Taylor series basis $\{d_1, d_2\}$, where we apply second-order Taylor series expansion on $V_{n+1}(x_{n+1})$ around $\hat{x}_{n+1} = A_n x_n + B_n a_n$. That is,

$$\begin{aligned} &\widetilde{M}(\alpha, \mathbf{z}) \\ &= \sum_{n=0}^{N-1} \left\{ \frac{\partial}{\partial x} V_{n+1}^T(\hat{x}_{n+1}) z_{n+1} + z_{n+1}^T \frac{\partial^2}{\partial x^2} V_{n+1}(\hat{x}_{n+1}) z_{n+1} - \mathbb{E} \left[z_{n+1}^T \frac{\partial^2}{\partial x^2} V_{n+1}(\hat{x}_{n+1}) z_{n+1} \right] \right\} \\ &= \sum_{n=0}^{N-1} \left\{ 2(A_n x_n + B_n a_n)^T K_{n+1} z_{n+1} + z_{n+1}^T K_{n+1} z_{n+1} - \mathbb{E} [z_{n+1}^T K_{n+1} z_{n+1}] \right\} = M^*(\alpha, \mathbf{z}). \end{aligned}$$

Therefore, the second-order Taylor series expansion of the optimal dual penalty is exact, which is due to the linearity of state dynamics and quadratic structure of the optimal value function. If using the approximation scheme in (4.4), we would have

$$\widehat{M}(\alpha, \mathbf{z}) = \sum_{n=0}^{N-1} \left(\frac{\partial}{\partial x} V_n(x_n) \right)^T z_{n+1} = \sum_{n=0}^{N-1} 2x_n^T K_n z_{n+1} \neq M^*(\alpha, \mathbf{z}).$$

5 Numerical Experiments: Dynamic Trading

5.1 Dynamic Trading with Predictable Returns and Transaction Costs

To demonstrate the effectiveness of the proposed regression approach in generating accurate approximations of optimal dual penalty and tight upper bounds, we will study the dynamic

trading problem with predictable returns and transaction costs in Gârleanu and Pedersen (2013). This model has been empirically tested in real financial markets and exhibits nice structural properties. The problem formulation is as follows.

Consider an investor that attempts to trade D securities at each time $t \in \{0, 1, \dots, T\}$. The securities' price changes (returns) are driven by K market factors, i.e.,

$$r_{t+1} = \mu_t + Bf_t + z_{t+1}^{(1)}, \quad t = 0, \dots, T-1,$$

where r_{t+1} is the $D \times 1$ vector of security returns at time $t+1$, μ_t is the deterministic “risk-free” return, f_t is the $K \times 1$ vector of market factors that predict returns, B is the $D \times K$ matrix of constant factor loadings, and $\{z_{t+1}^{(1)}\}$ are i.i.d. zero-mean $D \times 1$ random vectors with covariance matrix $Var(z_{t+1}^{(1)}) = \Sigma$ that represent the unpredictable noise in the return. Further assume the market return vector f_t is already known to the investor at time t (before the trading) and it follows a self-evolving state dynamics

$$\Delta f_t \triangleq f_{t+1} - f_t = -\Phi f_t + z_{t+1}^{(2)}, \quad t = 0, \dots, T-1, \quad (5.1)$$

where Φ is a $K \times K$ matrix of mean-reversion coefficients for the factors, and $\{z_{t+1}^{(2)}\}$ are i.i.d. $K \times 1$ zero-mean random vectors with covariance matrix $Var(z_{t+1}^{(2)}) = \Psi$ that represent the shock affecting the predictors. Usually, we assume $\{z_{t+1}^{(2)}\}$ are normally distributed. Assume Ψ satisfies the standard conditions such that f is stationary. Note that the state dynamics (5.1) could be rewritten as $f_{t+1} = (I - \Phi)f_t + z_{t+1}^{(2)}$, where I is the $K \times K$ identity matrix.

Assuming trading is costly in the sense that it impacts the market and results in a price move. In particular, Gârleanu and Pedersen (2013) argues that the transaction cost associated with trading quantity $a_t = x_t - x_{t-1}$, where x_t is the share quantity held by the investor at time t (after trading) and a_t is the trading quantity at time t , is given by $TC(a_t) = \frac{1}{2}a_t^T \Lambda a_t$. Here Λ is a $D \times D$ deterministic positive semi-definite matrix measuring the level of trading cost. The intuition is that trading a_t shares will move the price by $\frac{1}{2}\Lambda a_t$ and result in a total trading cost $(\frac{1}{2}\Lambda a_t)^T a_t$, which is exactly $TC(a_t)$.

The investor's objective is to choose a (non-anticipative) trading strategy that maximizes the total expected excess return with the trading costs and risk penalties taken into account, i.e.,

$$(P) : V_1(x_0, f_1) = \max_{\alpha \in \mathbb{A}_{\mathcal{F}}} \mathbb{E}_1 \left[\sum_{t=1}^T \left(x_t^T B f_t - \frac{1}{2} a_t^T \Lambda a_t - \frac{\gamma}{2} x_t^T \Sigma x_t \right) \right], \quad (5.2)$$

where $\mathbb{A}_{\mathcal{F}}$ describes the set of all admissible policies that satisfy certain trading constraints such as no short-sells or complete liquidation of the initial position by the end of trading horizon T , $\mathcal{F} = \{\mathcal{F}_0, \dots, \mathcal{F}_T\}$ is the natural information filtration, $x_t^T B f_t$ is the expected intermediate excess return, and γ is the risk-aversion coefficient.

Notice that the primal problem (5.2) falls into the realm of LQC if no trading constraints are imposed (except for complete liquidation of the initial position). Therefore, both the optimal value function and the optimal trading strategy could be solved in closed form using backward dynamic programming. For example, the optimal value function is quadratic in

the state (x_{t-1}, f_t) and the optimal control is an affine function of the state. In particular, for risk-neutral ($\gamma = 0$) trading, Moallemi and Saglam (2015) shows that the optimal value function $J_{t+1}(x_t, f_{t+1})$ satisfies

$$\begin{cases} J_T(x_{T-1}, f_T) = -\frac{1}{2}x_{T-1}^T \Lambda x_{T-1}, \\ J_t(x_{t-1}, f_t) = -\frac{1}{2}x_{t-1}^T A_{xx,t} x_{t-1} + x_{t-1}^T A_{xf,t} f_t + \frac{1}{2}f_t^T A_{ff,t} f_t + A_t, \quad t = T-1, \dots, 1, \end{cases} \quad (5.3)$$

where $A_{xx,t}$, $A_{xf,t}$, $A_{ff,t}$ and A_t are coefficient matrices of proper dimensions that satisfy the following backward recursions:

$$\begin{aligned} A_{xx,T} &= \Lambda, \quad A_{xf,T} = \mathbf{0}, \quad A_{ff,T} = \mathbf{0}, \quad A_T = 0; \\ A_{xx,t} &= -\Lambda(\Lambda + A_{xx,t+1})^{-1}\Lambda + \Lambda, \\ A_{xf,t} &= \Lambda(\Lambda + A_{xx,t+1})^{-1}(B + A_{xf,t+1}(I - \Phi)), \\ A_{ff,t} &= (B + A_{xf,t+1}(I - \Phi))^T(\Lambda + A_{xx,t+1})^{-1}(B + A_{xf,t+1}(I - \Phi)) + (I - \Phi)^T A_{ff,t+1}(I - \Phi), \\ A_t &= \text{trace}(\Psi A_{ff,t+1}) + A_{t+1}. \end{aligned}$$

Furthermore, the optimal policy α^* is given by

$$\alpha_t^*(x_{t-1}) = (\Lambda + A_{xx,t+1})^{-1}(\Lambda x_{t-1} + (B + A_{xf,t+1}(I - \Phi))f_t) - x_{t-1}. \quad (5.4)$$

Note that after imposing trading constraints, as many real-world trading problems do, the classic theory in LQC fails and in general it is difficult to find an optimal policy. Here let us consider a very natural set of trading constraints: only selling of securities is allowed, short-sell of the securities is not allowed, and the initial position must be completely liquidated by terminal horizon. Mathematically, the constraint set \mathcal{A} could be formulated as

$$\mathcal{A} \triangleq \{ \mathbf{a} = (a_1, \dots, a_T) : x_t = x_{t-1} + a_t, a_t \leq 0, x_t \geq 0, t = 1, \dots, T; x_T = 0 \}$$

Notice that all the constraints in \mathcal{A} are linear (hence convex) in actions. Thus, the problem maintains to be tractable when deterministic convex objective functions are of interest.

In the existence of trading constraints \mathcal{A} , the optimal trading strategy is not readily available. Alternatively, we focus on finding sufficiently good suboptimal strategies. Although the optimal policy of the unconstrained problem is unlikely to be feasible for the constrained problem (since one of the constraints in \mathcal{A} might be violated), a natural feasible policy could be derived by projecting the optimal policy of the unconstrained problem onto the constraint set \mathcal{A} . In particular, we consider the following projected linear quadratic control (PLQC) policy $\hat{\alpha} = (\hat{\alpha}_1, \dots, \hat{\alpha}_T)$ with

$$\hat{\alpha}_t(x_{t-1}) = \max\{-x_{t-1}, \min\{0, \alpha_t^*(x_{t-1})\}\}, \quad t = 1, \dots, T, \quad (5.5)$$

where the projection is componentwise and $\alpha_t^*(x_{t-1})$ is given by (5.4). From (5.5), the projection of control onto \mathcal{A} is easy to implement, and thus it is straightforward to evaluate the PLQC policy via Monte Carlo simulation. Eventually, it results in a lower bound $V_1^{\hat{\alpha}}(x_0, f_1)$ on the optimal value function $V_1(x_0, f_1)$.

Remark 5.1. In Moallemi and Saglam (2015), the authors also consider a time-weighted average price (TWAP) policy where an equal quantity of shares ($a_t = -x_0/T$) is traded

at every period, a deterministic (DETER) policy where a deterministic trading strategy is given by solving the deterministic trading problem with all noises ignored, a model predictive control (MPC) policy where the trading strategy is given by solving a sequence of deterministic trading problems without noises, and a linear rebalancing (LRB) policy where the control is given by the best rebalancing policy that is affine in all the available market predictors. Our preliminary numerical tests show that these three policies (TWAP, DETER, MPC) perform significantly worse than the PLQC policy. Although the LRB policy is competitive (often worse) compared with the PLQC policy in terms of performance; it is more sensitive to formulation, more complex in structure, and more difficult to solve. Therefore, we focus on the easy-to-implement PLQC policy.

5.2 Dual Formulation with Regression-based Penalties

Now let us consider the dual formulation of primal problem (5.2). Following the derivation in Section 2.1, we know that the dual problem

$$(D) : V_1^M(x_0, f_1) = \mathbb{E}_1 \left[\max_{\alpha \in \mathbb{A}} \sum_{t=1}^T \left(x_t^T B f_t - \frac{1}{2} a_t^T \Lambda a_t - \frac{\gamma}{2} x_t^T \Sigma x_t \right) - M(\alpha, \mathbf{z}^{(2)}) \right], \quad (5.6)$$

where $\mathbf{z}^{(2)} := (z_1^{(2)}, \dots, z_T^{(2)})$ and $M(\alpha, \mathbf{z}^{(2)})$ is a feasible dual penalty. It follows that $V_1^M(x_0, f_1) \geq V_1(x_0, f_1)$. In particular, when plugging in the optimal dual penalty $M^*(\alpha, \mathbf{z}^{(2)})$ the exact value function is recovered, i.e., $V_1^{M^*}(x_0, f_1) = V_1(x_0, f_1)$, where

$$M^*(\alpha, \mathbf{z}^{(2)}) = \sum_{t=1}^{T-1} \left(V_{t+1}(x_t, f_{t+1}) - \mathbb{E}[V_{t+1}(x_t, f_{t+1}) | x_{t-1}, a_t, f_t] \right).$$

Since the optimal value functions $\{V_{t+1}(x_t, f_{t+1})\}$ are not available, we have to approximate $M^*(\alpha, \mathbf{z}^{(2)})$. In particular, we will apply the proposed regression approach. Note that the objective function is quadratic, we will use the second-order Taylor series basis (which is equivalent to the second-order L^2 orthonormal basis for this problem) in the algorithm.

Following the derivation in Section 3.3, let us express the optimal dual penalty w.r.t. the basis $\{d_1 = z, d_2 = z^2 - \mathbb{E}[z^2]\}$, i.e.,

$$M^*(\alpha, \mathbf{z}^{(2)}) \approx \sum_{t=1}^{T-1} \left(\gamma_{t,1}^T(x_{t-1}, a_t, f_t) z_{t+1}^{(2)} + \gamma_{t,2}^T(x_{t-1}, a_t, f_t) \left((z_{t+1}^{(2)})^2 - \mathbb{E}[(z_{t+1}^{(2)})^2] \right) \right),$$

where $\{\gamma_{t,r}(x_{t-1}, a_t, f_t)\}$ are computed via

$$\begin{cases} \gamma_{t,1}(x_{t-1}, a_t, f_t) = \text{Cov}(z_{t+1}^{(2)})^{-1} \mathbb{E}_t \left[V_{t+1}(x_t, f_{t+1}) z_{t+1}^{(2)} \right] \\ \gamma_{t,2}(x_{t-1}, a_t, f_t) = \text{Cov}((z_{t+1}^{(2)})^2)^{-1} \mathbb{E}_t \left[V_{t+1}(x_t, f_{t+1}) \left((z_{t+1}^{(2)})^2 - \mathbb{E}[(z_{t+1}^{(2)})^2] \right) \right]. \end{cases} \quad (5.7)$$

Note that we are slightly abusing the notation $(z_{t+1}^{(2)})^2$ to denote the componentwise square of a random vector. We will use Algorithm 1 to estimate $\{\gamma_{t,r}(x_{t-1}, a_t, f_t)\}$ in (5.7). The key is how to construct proper regressors in the regression. As discussed in the previous section, a good candidate regressor is $\frac{\partial r}{\partial f^r} \tilde{V}_{t+1}(x_t, \tilde{f}_{t+1}) \cdot (\sigma)^r$, where $\tilde{V}_{t+1}(\cdot)$ is either the value function of

a naive policy or the value function of a simpler DP such that its derivatives admit closed-form expressions. Here $\hat{f}_{t+1} = (I - \Phi)f_t$. Recall that the corresponding unconstrained problem is an LQC problem, and its value functions $\{J_{t+1}(x_t, f_{t+1})\}$ admit closed-form expressions as in (5.3). Therefore, we can use $\frac{\partial^r}{\partial f^r} J_{t+1}(x_t, \hat{f}_{t+1}) \cdot (\sigma)^r$ as one of the regressors for regressing $\gamma_{n,r}(x_{t-1}, a_t, f_t)$. Moreover, $J_{t+1}(x_t, \hat{f}_{t+1})$ is quadratic in (x_t, \hat{f}_{t+1}) , x_t is affine in (x_{t-1}, a_t) , and \hat{f}_{t+1} is linear in f_t . Thus, the first-order and second-order derivatives of $J_{t+1}(x_t, \hat{f}_{t+1})$ w.r.t. f are both affine in (x_{t-1}, a_t, f_t) . It follows that the regressor, and thus the resulted dual penalty is affine in policy and preserves convexity in the dual problem. In particular, our numerical tests show that the regressors

$$\phi_{t,1}(x_{t-1}, a_t, f_t) = \left\{ 1, \frac{\partial}{\partial f} J_{t+1}(x_t, \hat{f}_{t+1}) = A_{xf,t+1}^T(x_{t-1} + a_t) + A_{ff,t+1}(I - \Phi)f_t \right\}$$

and

$$\phi_{t,2} = \left\{ \frac{\partial^2}{\partial f^2} J_{t+1}(x_t, \hat{f}_{t+1}) = \text{diag}(A_{ff,t+1}) \right\} \quad (5.8)$$

are sufficiently good for regressing $\gamma_{n,1}(x_{t-1}, a_t, f_t)$ and $\gamma_{n,2}(x_{t-1}, a_t, f_t)$, respectively. Here note that the derivatives are componentwise, $\text{diag}(A_{ff,t+1})$ denotes the diagonal vector of matrix $A_{ff,t+1}$, and we suppress the dependence of function $\phi_{t,2}$ on (x_{t-1}, a_t, f_t) because $\text{diag}(A_{ff,t+1})$ is a constant. To this end, let us denote the resulted regression-based dual penalty as

$$\widetilde{M}^t(\alpha, \mathbf{z}^{(2)}) = \sum_{t=1}^{T-1} \left(\widetilde{\gamma}_{t,1}^T(x_{t-1}, a_t, f_t) z_{t+1}^{(2)} \widetilde{\gamma}_{t,2}^T(x_{t-1}, a_t, f_t) \left((z_{t+1}^{(2)})^2 - \mathbb{E} \left[(z_{t+1}^{(2)})^2 \right] \right) \right), \quad (5.9)$$

where $\widetilde{\gamma}_{t,r}(x_{t-1}, a_t, f_t)$ is the estimation of $\gamma_{t,r}(x_{t-1}, a_t, f_t)$ after regression. Notice that the regressor $\phi_{t,2}$ in (5.8) does not depend on (x_{t-1}, a_t, f_t) , it follows that $\widetilde{\gamma}_{t,2}^T$ does not depend on (x_{t-1}, a_t, f_t) . Therefore, the second component of $\widetilde{M}^t(\alpha, \mathbf{z}^{(2)})$,

$$\sum_{t=1}^{T-1} \widetilde{\gamma}_{t,2}^T \left((z_{t+1}^{(2)})^2 - \mathbb{E} \left[(z_{t+1}^{(2)})^2 \right] \right)$$

does not involve states and actions. Thus, it does not affect the optimization in the inner optimization problems. In fact, it plays the role of control variate in the sense that it has zero mean and thus does not affect the expectation of the upper bound estimator; however, the correlation between this term and the reward function might help reduce its variance. Plugging $\widetilde{M}^t(\alpha, \mathbf{z}^{(2)})$ into the dual problem (5.6), we have

$$V_1^{\widetilde{M}^t}(x_0, f_1) = \mathbb{E}_1 \left[\max_{\alpha \in \mathbb{A}} \sum_{t=1}^T \left(x_t^T B f_t - \frac{1}{2} a_t^T \Lambda a_t - \frac{\gamma}{2} x_t^T \Sigma x_t \right) - \widetilde{M}^t(\alpha, \mathbf{z}^{(2)}) \right].$$

Note that $\widetilde{M}^t(\alpha, \mathbf{z}^{(2)})$ is a feasible dual penalty, and thus $V_1^{\widetilde{M}^t}(x_0, f_1) \geq V_1(x_0, f_1)$. To obtain an unbiased estimator of $V_1^{\widetilde{M}^t}(x_0, f_1)$, we can simulate L i.i.d. sample paths of $\mathbf{z}^{(2)}$, solve the deterministic inner optimization problem corresponding to every sample path, and finally take the sample average of the optimal values. Since the inner optimization problems are deterministic and convex, we use the CVX package (Mosek solver) in Matlab to solve them.

5.3 Numerical Results

We will use the model parameters calibrated in Moallemi and Saglam (2015). However, the authors only consider the case of trading one stock (Apple, Inc), i.e., $D = 1$. To exhibit the effectiveness of our regression approach for high-dimensional DPs, we “replicate” their model parameters to obtain a high-dimensional dynamic trading problem. The model parameters are summarized as follows.

Time horizon is $T = 12$ or 24 . The number of stocks to be liquidated is $D = 5, 10$, or 25 . The initial position of every stock to be liquidated is $[x_0^1, \dots, x_0^D]$, a vector in \mathbb{R}^D with $x_0^d = 10000$. There are $K = 2$ market factors to predict the return on each stock, each with a different mean reversion speed. The return of the stocks follows the dynamics

$$r_{t+1}^d = 0.0726 + 0.3375f_{t,1} - 0.0720f_{t,2} + z_{t+1,d}^{(1)}, \quad d = 1, \dots, D; \quad t = 0, \dots, T-1,$$

where $\Sigma = \text{Var}(z_{t+1,d}^{(1)}) = 0.048$. Thus, the $D \times K$ matrix

$$B = [0.3375, -0.072; 0.3375, -0.072; \dots, 0.3375, -0.072].$$

Similarly, the market factor f follows the dynamics

$$\begin{cases} f_{t+1,1} - f_{t,1} = -0.5f_{t,1} + z_{t+1,1}^{(2)}, \\ f_{t+1,2} - f_{t,2} = -0.7f_{t,2} + z_{t+1,2}^{(2)}, \end{cases}$$

where

$$\Psi = \text{Var}(z_{t+1}^{(2)}) = \begin{bmatrix} 0.0379 & 0 \\ 0 & 0.0947 \end{bmatrix},$$

and

$$\Phi = \begin{bmatrix} 0.5 & 0 \\ 0 & 0.7 \end{bmatrix}.$$

We also test other choices of Φ in the appendix. Instead of generating f_0, f_1 randomly, we let $f_0 = [1, 1]'$ and $f_1 = (I - \Phi)f_0$. The transaction cost matrix $\Lambda = \lambda \cdot \tilde{\Lambda}$, where $\lambda = 2.14 \times 10^{-5}$, and $\tilde{\Lambda} = \Gamma \Gamma^T$. Here Γ is an upper triangular matrix such that

$$\Gamma = \begin{bmatrix} \frac{1}{\sqrt{D}} & \frac{1}{\sqrt{D}} & \cdots & \frac{1}{\sqrt{D}} \\ 0 & \frac{1}{\sqrt{D-1}} & \cdots & \frac{1}{\sqrt{D-1}} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 1 \end{bmatrix}_{D \times D}.$$

The reason for using such a matrix is to ensure that diagonal elements of $\tilde{\Lambda}$ are all one, meaning the transaction costs are the same across all the stocks. We also test other choices of λ in the appendix.

With the above model parameters, we solve for the parameters (i.e., A_{xx} , A_{xf} , A_{ff} , and A) of the optimal policy to the corresponding unconstrained problem. Then the PLQC policy is computed via policy projection (5.4). To evaluate the PLQC policy and generate a lower bound estimator, we run a simulation of $M = 10^6$ i.i.d. sample paths with the same

initial state (x_0, f_1) , exercise the PLQC policy along each sample path, and compute the corresponding accumulated reward. The value function of the PLQC policy, i.e., the lower bound on the optimal value function, is estimated by taking the average of those reward samples. As for the upper bound estimation, we consider the upper bounds induced by three different feasible dual penalties: (1) Zero penalty (note that it is a trivial feasible dual penalty), meaning the investor has a perfect foresight of all the future information without any penalization; (2) Regression-based penalty with the first-order Taylor series basis, i.e.,

$$\widetilde{M}_1^t(\boldsymbol{\alpha}, \mathbf{z}^{(2)}) = \sum_{t=1}^{T-1} \widetilde{\gamma}_{t,1}^T(x_{t-1}, a_t, f_t) z_{t+1}^{(2)}; \quad (5.10)$$

(3) Regression-based penalty with the second-order Taylor series basis, i.e., $\widetilde{M}^t(\boldsymbol{\alpha}, \mathbf{z}^{(2)})$ in (5.9). To regress $\{\gamma_{t,r}^T(x_{t-1}, a_t, f_t)\}$, we implement Algorithm 1 with the PLQC policy and the $M = 10^6$ sample paths generated in estimating the lower bound. After constructing all three dual penalties, we simulate $L = 400$ sample paths to generate the upper bound estimators. L is much smaller than M because the variances of the upper bound estimators induced by good approximations of the optimal dual penalty are small. To illustrate the performance of the PLQC policy, we compute its duality gap, which is the ratio of the difference between the lower bound and the tightest upper bound to the lower bound. If the duality gap is small enough, we could claim the policy is sufficiently good. The detailed results are summarized in Table 5.1.

Table 5.1: Dual Bounds of PLQC Policy with Parameter Φ .

D	T	LB PLQC (\$k\$)	UB 1 ZP (\$k\$)	UB 2 F-O (\$k\$)	UB 3 S-O (\$k\$)	DualGap
5	12	14.937 (0.074)	18.096 (1.304)	15.365 (0.107)	15.263 (0.055)	2.18%
10	12	24.303 (0.140)	30.240 (2.224)	24.544 (0.152)	24.645 (0.070)	0.99%
25	12	32.090 (0.324)	44.124 (5.237)	32.657 (0.211)	32.529 (0.109)	1.37%
5	24	18.635 (0.103)	28.060 (2.094)	20.012 (0.250)	20.103 (0.152)	7.39%
10	24	33.309 (0.212)	47.738 (3.816)	35.099 (0.306)	35.285 (0.193)	5.37%
25	24	62.971 (0.528)	86.537 (9.673)	64.401 (0.370)	65.224 (0.201)	2.27%

The experiment parameters: $M = 10^6$ sample paths are simulated to evaluate the lower bounds and compute the coefficients in the regression; $L = 400$ sample paths are simulated to evaluate the upper bounds.

In Table 5.1, Column 1 records the number of securities D , including the cases where $D = 5, 10$ or 25 . Column 2 records the number of trading horizons T , including the cases $T = 12$ or 24 . Column 3 records the lower bounds (LB) induced by the PLQC policy. Column 4 records the upper bounds (UB 1) induced by the zero penalty. Column 5 records the upper

bounds (UB 2) induced by the dual penalty $\widetilde{M}_1^t(\boldsymbol{\alpha}, \mathbf{z}^{(2)})$ in (5.10). Column 6 records the upper bounds (UB 3) induced by the dual penalty $\widetilde{M}^t(\boldsymbol{\alpha}, \mathbf{z}^{(2)})$ in (5.9). The half confidence interval widths of the lower bound and each upper bound are presented in the parentheses. The last column records the best duality gaps (DualGap) achieved by comparing the lower bounds and the tightest upper bounds. We have the following observations:

- In general, in view of the small duality gaps in most cases, we could claim that the simple PLQC policy is a sufficiently good policy.
- Comparing UB 1, UB 2 and UB 3 induced by the zero penalty, the first-order regression-based penalty, and the second-order regression-based penalty, respectively, we notice: 1) UB 1 has large duality gaps, meaning the zero penalty performs poorly as one would expect. 2) For most parameter setups, UB 2 and UB 3 have very small duality gaps, meaning that the policy used in the regression (here is the PLQC policy) is near optimal and the regression-based dual penalties are accurate approximations of the optimal dual penalty. 3) It is hard to distinguish between UB 2 and UB 3 in terms of tightness. The reason is that the inner optimization problems are the same, except for a control variate term in the second-order regression-based penalty. This term helps reduce the variance of the upper bound estimator, which is verified by the fact that the half confidence interval widths of UB 3 are narrower than the ones of UB 2.
- We also observe that, as the number of trading horizons T increases, the duality gaps increase. It indicates the possibility that the PLQC policy becomes less optimal for larger T . Another possibility is that the approximation of the optimal dual penalty is less accurate for larger T .
- Lastly, as the number of securities in position increases, the duality gaps decrease instead of increasing as one would expect. One possible explanation is that the PLQC policy is closer to optimal for higher-dimensional problems due to the strong correlations across the assets in the transaction cost matrix Λ . Therefore, the optimal trading strategy of the unconstrained problem is more conservative, and thus less optimality is lost in the policy projection.

Overall, we can see that the PLQC policy is near optimal for the model parameters tested. More importantly, the regression-based dual penalties perform well. It demonstrates the effectiveness and efficiency of the proposed framework of regression approach in solving the dual problems of high-dimensional DPs.

6 Conclusions

In this paper, we develop a framework of regression approach to approximating the optimal dual penalty in general DPs, by studying the structure of the dual penalty space. The proposed framework circumvents the issue of nested simulation suffered by some of the existing approaches, and thus improves computational efficiency. Furthermore, the proposed framework requires minimal extra simulation and computational costs by reusing the samples generated during the estimation of lower bound. Lastly, the proposed framework is robust. Dual penalties with desired structural properties could be generated by constructing proper regressors in the regression. The application to a high-dimensional dynamic trading problem

demonstrates its effectiveness in generating good feasible dual penalties and tight upper bounds on the optimal value function.

References

- Anderson, L., M. Broadie. 2004. Primal-dual simulation algorithm for pricing multidimensional American options. *Management Science* **50**(9) 1222 – 1234.
- Belomestny, D., C. Bender, J. Schoenmakers. 2009. True upper bounds for Bermudan products via non-nested Monte Carlo. *Mathematical Finance* **19** 53 – 71.
- Bertsekas, D.P. 2007. *Dynamic Programming and Optimal Control*. 3rd ed. Athena Scientific.
- Bogachev, Vladimir I. 2007. *Measure Theory*, vol. 1. Springer Science & Business Media.
- Brown, D., J. Smith, P. Sun. 2010. Information relaxations and duality in stochastic dynamic programs. *Operations Research* **58**(4) 758 – 801.
- Brown, David B, James E Smith. 2014. Information relaxations, duality, and convex stochastic dynamic programs. *Operations Research* **62**(6) 1394–1415.
- Brown, D.B., J.E. Smith. 2011. Dynamic portfolio optimization with transaction costs: Heuristics and dual bounds. *Management Science* **57**(10) 1752–1770.
- Chang, H. S., M. C. Fu, J. Hu, S. I. Marcus. 2007. *Simulation-based Algorithms for Markov Decision Processes*. 1st ed. Communications and Control Engineering Series, Springer, New York.
- Chen, Nan, Paul Glasserman. 2007. Additive and multiplicative duals for american option pricing. *Finance and Stochastics* **11**(2) 153–179.
- Desai, V., V. Farias, C. Moallemi. 2011. Bounds for Markov decision processes. Chapter in Reinforcement Learning and Approximate Dynamic Programming for Feedback Control (F. L. Lewis, D. Liu, eds.).
- Gârleanu, Nicolae, Lasse Heje Pedersen. 2013. Dynamic trading with predictable returns and transaction costs. *The Journal of Finance* **68**(6) 2309–2340.
- Haugh, M., L. Kogan. 2004. Pricing American options: A duality approach. *Operations Research* **52**(2) 258 – 270.
- Haugh, Martin, Andrew EB Lim. 2012. Linear–quadratic control and information relaxations. *Operations Research Letters* **40**(6) 521–528.
- Kim, Michael J., Andrew EB Lim. 2015. Robust multiarmed bandit problems. *Management Science* **62**(1) 264–285.
- Lai, G., F. Margot, N. Secomandi. 2010. An approximate dynamic programming approach to benchmark practice-based heuristics for natural gas storage valuation. *Operations Research* **58**(3) 564–582.
- Moallemi, Ciamac C, Mehmet Saglam. 2015. Dynamic portfolio choice with linear rebalancing rules. *Available at SSRN 2011605* .
- Powell, W.B. 2011. *Approximate Dynamic Programming: Solving the curses of dimensionality*. 2nd ed. John Wiley and Sons.
- Rogers, L. 2002. Monte Carlo valuation of American options. *Mathematical Finance* **12**(3) 271 – 286.
- Rogers, LCG. 2007. Pathwise stochastic optimal control. *SIAM Journal on Control and Optimization* **46**(3) 1116–1132.

- Ye, F., E. Zhou. 2012. Parameterized penalties in the dual representation of markov decision processes. *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*. IEEE, 870–876.
- Ye, Fan, Enlu Zhou. 2015. Information relaxation and dual formulation of controlled markov diffusions. *Automatic Control, IEEE Transactions on* **60**(10) 2676–2691.
- Zhu, Helin, Fan Ye, Enlu Zhou. 2015. Fast estimation of true bounds on bermudan option prices under jump-diffusion processes. *Quantitative Finance* **15**(11) 1885–1900.

Appendix

The performances of the regression-based penalties under different choices of factor persistence Φ . In particular,

$$\Phi_1 = \begin{bmatrix} 0.3 & 0 \\ 0 & 0.5 \end{bmatrix}, \Phi_2 = \begin{bmatrix} 0.3 & 0 \\ 0 & 0.7 \end{bmatrix}, \Phi_3 = \begin{bmatrix} 0.5 & 0 \\ 0 & 0.3 \end{bmatrix}, \Phi_4 = \begin{bmatrix} 0.7 & 0 \\ 0 & 0.5 \end{bmatrix}.$$

Table 6.1: Lower and Upper Bounds of PLQC policy.

Φ	D	T	Lower Bound PLQC Policy (\$k\$)	UpperBound 1 Zero Penalty (\$k\$)	UpperBound 2 First-Order (\$k\$)	UpperBound 3 Second-Order (\$k\$)	Duality Gap
Φ_1	1	12	7.564 (0.026)	8.327 (0.448)	7.760 (0.070)	7.732 (0.046)	2.22%
Φ_1	5	12	33.571 (0.122)	36.675 (1.890)	33.692 (0.191)	33.670 (0.103)	0.29%
Φ_1	10	12	59.391 (0.230)	63.451 (3.509)	59.893 (0.291)	59.837 (0.136)	0.75%
Φ_1	25	12	110.277 (0.506)	131.213 (8.102)	112.164 (0.552)	113.531 (0.326)	1.71%
Φ_2	1	12	7.964 (0.026)	9.311 (0.815)	8.155 (0.066)	8.170 (0.042)	2.40%
Φ_2	5	12	35.460 (0.122)	39.136 (1.787)	35.955 (0.241)	35.758 (0.096)	0.84%
Φ_2	10	12	63.287 (0.228)	70.992 (3.658)	63.770 (0.261)	63.885 (0.117)	0.76%
Φ_2	25	12	119.974 (0.505)	135.117 (9.307)	123.230 (0.526)	122.642 (0.277)	2.22%
Φ_3	1	12	2.565 (0.014)	3.917 (0.222)	2.791 (0.063)	2.849 (0.049)	8.81%
Φ_3	5	12	9.776 (0.070)	13.097 (1.217)	10.072 (0.150)	10.374 (0.084)	3.02%
Φ_3	10	12	13.946 (0.135)	20.531 (2.373)	14.388 (0.193)	14.489 (0.107)	3.17%
Φ_3	25	12	6.897 (0.327)	14.942 (5.140)	7.392 (0.250)	7.372 (0.142)	6.89%
Φ_4	1	12	1.274 (0.010)	2.029 (0.166)	1.463 (0.040)	1.445 (0.026)	13.42%
Φ_4	5	12	3.609 (0.051)	5.978 (0.817)	3.715 (0.081)	3.756 (0.043)	2.94%
Φ_4	10	12	2.638 (0.099)	8.173 (1.657)	2.741 (0.162)	2.707 (0.064)	2.62%
Φ_4	25	12	-18.907 (0.245)	-5.989 (3.906)	-18.356 (0.159)	-18.448 (0.081)	2.43%

The experiment parameters: $M = 10^6$ sample paths are simulated to evaluate the lower bounds and generate the regression-based penalties; $L = 400$ sample paths are simulated to generate the upper bounds.

The performances of the regression-based penalties with different choices of transaction cost parameter λ . In particular, we let

$$\lambda_1 = 1.07 \times 10^{-5}, \lambda_2 = 2.67 \times 10^{-5}, \lambda_3 = 3.21 \times 10^{-5}, \lambda_4 = 4.28 \times 10^{-5}.$$

Table 6.2: Lower and Upper Bounds of PLQC policy.

λ	D	T	Lower Bound PLQC Policy (\$k\$)	UpperBound 1 Zero Penalty (\$k\$)	UpperBound 2 First-Order (\$k\$)	UpperBound 3 Second-Order (\$k\$)	Duality Gap
λ_1	1	12	3.757 (0.015)	4.691 (0.250)	4.104 (0.078)	4.027 (0.056)	7.19%
λ_1	5	12	16.963 (0.075)	21.195 (1.324)	17.210 (0.157)	17.353 (0.094)	1.46%
λ_1	10	12	29.899 (0.145)	36.332 (2.444)	30.377 (0.226)	29.749 (0.104)	1.60%
λ_1	25	12	55.321 (0.338)	65.456 (6.056)	56.405 (0.343)	55.902 (0.201)	1.05%
λ_2	1	12	3.583 (0.015)	4.286 (0.234)	3.759 (0.045)	3.756 (0.034)	4.83%
λ_2	5	12	14.119 (0.072)	18.721 (1.221)	14.437 (0.090)	14.433 (0.040)	2.25%
λ_2	10	12	21.854 (0.136)	25.148 (2.096)	22.444 (0.120)	22.458 (0.062)	2.70%
λ_2	25	12	21.162 (0.320)	28.615 (5.446)	21.743 (0.200)	21.862 (0.117)	2.75%
λ_3	1	12	3.533 (0.015)	4.552 (0.251)	3.626 (0.039)	3.642 (0.023)	2.63%
λ_3	5	12	13.326 (0.072)	18.707 (1.083)	13.694 (0.096)	13.649 (0.039)	2.42%
λ_3	10	12	19.759 (0.134)	26.425 (2.290)	20.091 (0.099)	20.254 (0.045)	1.68%
λ_3	25	12	10.798 (0.318)	20.543 (4.439)	11.275 (0.161)	11.164 (0.085)	3.39%
λ_4	1	12	3.417 (0.015)	4.486 (0.273)	3.523 (0.035)	3.510 (0.018)	2.72%
λ_4	5	12	12.008 (0.070)	16.255 (1.120)	12.275 (0.079)	12.241 (0.031)	1.94%
λ_4	10	12	15.811 (0.132)	19.592 (2.070)	16.215 (0.092)	16.182 (0.043)	2.35%
λ_4	25	12	-9.375 (0.315)	3.890 (4.766)	-9.258 (0.155)	-9.335 (0.108)	0.43%

The experiment parameters: $M = 10^6$ sample paths are simulated to evaluate the lower bounds and generate the regression-based penalties; $L = 400$ sample paths are simulated to generate the upper bounds.